

Detecting Human-Object Interactions with Object-Guided Cross-Modal Calibrated Semantics

Hangjie Yuan¹, Mang Wang³, Dong Ni^{1,2*}, Liangpeng Xu³

¹ College of Control Science and Engineering, Zhejiang University, Hangzhou, China

² State Key Laboratory of Industrial Control Technology, Zhejiang University, Hangzhou, China

³ DAMO Academy, Alibaba Group, China

{hj.yuan, dni}@zju.edu.cn, {wangmang.wm, liangpeng.xlp}@alibaba-inc.com

Abstract

Human-Object Interaction (HOI) detection is an essential task to understand human-centric images from a fine-grained perspective. Although end-to-end HOI detection models thrive, their paradigm of parallel human/object detection and verb class prediction loses two-stage methods' merit: object-guided hierarchy. The object in one HOI triplet gives direct clues to the verb to be predicted. In this paper, we aim to boost end-to-end models with object-guided statistical priors. Specifically, We propose to utilize a Verb Semantic Model (VSM) and use semantic aggregation to profit from this object-guided hierarchy. Similarity KL (SKL) loss is proposed to optimize VSM to align with the HOI dataset's priors. To overcome the static semantic embedding problem, we propose to generate cross-modality-aware visual and semantic features by Cross-Modal Calibration (CMC). The above modules combined composes Object-guided Cross-modal Calibration Network (OCN). Experiments conducted on two popular HOI detection benchmarks demonstrate the significance of incorporating the statistical prior knowledge and produce state-of-the-art performances. More detailed analysis indicates proposed modules serve as a stronger verb predictor and a more superior method of utilizing prior knowledge. The codes are available at <https://github.com/JacobYuan7/OCN-HOI-Benchmark>.

Introduction

Human Object Interaction (HOI) detection has recently become a thriving research topic as it provides a fine-grained understanding to human-centric images. HOI detection aims to detect triplets formulated as $\langle \text{human}, \text{verbs}, \text{object} \rangle$. Elaborate HOI detection can boost the results of image captioning (Yao et al. 2018), image retrieval (Johnson et al. 2015), activity recognition (Yuan, Ni, and Wang 2021), etc..

Recently proposed end-to-end methods (Liao et al. 2020; Kim et al. 2021; Tamura, Ohashi, and Yoshinaga 2021) for HOI detection have achieved notable results without two-stage processing. Unlike two-stage methods that deal with human/object detection and verb prediction sequentially, end-to-end methods run these two processes in parallel. Since two-stage methods identify object class first, it is more effective than end-to-end methods in leveraging the

object class information for verb prediction. Objects can reveal direct clues to the interactiveness of HOI triplets (Li et al. 2019) and furthermore the specific interactions of HOI triplets. Although those unlikely object-verb pairs can be rejected by introducing a post-processing (e.g. applying a binary mask (Tamura, Ohashi, and Yoshinaga 2021)), a better choice is to inject statistical prior information during training, thus producing better ranking results (Chen et al. 2019). In this paper, we take a step further to implant the prior knowledge into an end-to-end model.

To profit from the object-guided hierarchy, we propose to utilize the semantic space (Rahman et al. 2020; Xu et al. 2019). Semantic embeddings are initialized from word embeddings (Pennington, Socher, and Manning 2014; Mikolov et al. 2013). Several attempts on utilizing semantic space in HOI constrained in vanilla pretrained embeddings or their independent projections (i.e. MLPs) (Zhong et al. 2020; Gao et al. 2020; Bansal et al. 2020; Peyre et al. 2019). However, semantic space tends to have a domain discrepancy towards visual space (Zhu et al. 2021), which can not be overcome by such transformations. Some works (Xu et al. 2019; Peyre et al. 2019) incorporate semantic space to push visual space closer to it, which will cause negative effects when visual features are strong and robust. Inspired by (Wu et al. 2018; You et al. 2020), we propose a Verb Semantic Model (VSM) that outputs a set of semantic features that fit in with the HOI dataset's verb co-occurrence priors. This procedure is optimized by the proposed Similarity KL (SKL) loss without entry relaxation. Then, we inject object-verb hierarchical priors via semantic aggregation, which generates a set of semantic features corresponding to a set of visual features. By proper utilization of two-modal features, performance can still be boosted under strong vision models.

As VSM is shared across all images for a given dataset, semantic aggregation gathers static verb semantic embeddings given an object class. This static property is identical to previous methods (Liu, Chen, and Zisserman 2020) and lacks cross-modality-aware representation. As two modalities provide complementarity, we manage to solve the problem by cross-modal one-to-one calibration: mutually calibrating features from one modality via excitation from the other modality. By this mutual calibration, our model can generate a vision-aware semantic feature set and a semantic-aware visual feature set. Specifically, we propose to do

*Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Cross-Modal Calibration (CMC), which includes **i**) calibrating each modality’s features by the other modality via Inter-modal Calibration (InterC), **ii**) furthermore utilizing Intra-modal Enhanced Calibration (IntraEC) (Vaswani et al. 2017; Lin et al. 2020) to achieve intra-modal global reasoning.

In order to apply the proposed VSM and CMC into practice, we select an end-to-end vision model (VM) (Zou et al. 2021; Tamura, Ohashi, and Yoshinaga 2021) as our VM and compose Object-guided Cross-modal Calibration Network (OCN). To conclude, our contributions are three-fold:

- We introduce the object-guided statistical priors to facilitate end-to-end HOI detection. We introduce a Verb Semantic Model and use semantic aggregation to profit from this object-guided hierarchy. SKL loss is proposed to optimize VSM to align with the HOI dataset’s priors.
- To overcome the problem of static semantic embeddings, we propose to generate cross-modality-aware visual and semantic features by Cross-Modal Calibration, which consists of Inter-modal Calibration and Intra-modal Enhanced Calibration.
- Equipped with proposed modules, our end-to-end HOI detection model OCN achieves state-of-the-art results on two popular benchmarks. More detailed analysis indicates proposed modules serve as a stronger verb predictor and a more superior method of utilizing prior knowledge.

Related Work

Human-Object Interaction Detection Since the proposition of visual semantic role labeling and V-COCO dataset (Gupta and Malik 2015), HOI detection has been a research hotspot. The methods to tackle with this problem can be categorized into two-stage methods and one-stage methods.

Two-stage methods follow a pipeline of detecting objects first and then use cropped features to infer one HOI triplet’s multi-label verb interactions. The inference usually contain multiple streams (Chao et al. 2018): an object stream, a human stream and an interaction stream, which is the key stream. InteractNet (Gkioxari et al. 2018) is propose to predict action-specific density maps. Contextual attention (Wang et al. 2019a) is proposed to select contextual interaction information. IDN (Li et al. 2020) is proposed to learn interactions by pair integration and decomposition. Various graph-based models like GPNN (Qi et al. 2018), RPNN (Zhou and Chi 2019), VSGNet (Ulutan, Iftekhhar, and Manjunath 2020), DRG (Gao et al. 2020), CHG (Wang, Zheng, and Yingbiao 2020) are proposed to capture the interaction pattern from different aspects. Other cues like poses (Gupta, Schwing, and Hoiem 2019; Li et al. 2019), spatial layouts (Gao, Zou, and Huang 2018), action co-occurrence (Kim et al. 2020b) and language features (Liu, Chen, and Zisserman 2020; Xu et al. 2019; Peyre et al. 2019; Zhong et al. 2020) are utilized to augment HOI detection.

One-stage methods can be categorized into three types: **i**) point-based methods (Liao et al. 2020; Zhong et al. 2021) which infers heuristically-defined interaction points, **ii**) anchor-based methods (Kim et al. 2020a) which detects union boxes, **iii**) DETR-based (Carion et al. 2020) methods. Thanks to DETR and its ability to extract contextual cues

(Vaswani et al. 2017), several customized end-to-end models (Kim et al. 2021; Zou et al. 2021; Chen et al. 2021; Tamura, Ohashi, and Yoshinaga 2021) originated from DETR have achieved promising results. However, there has not been an end-to-end method to incorporate semantics or explicitly utilize object-guided hierarchical relation priors.

Language Semantics for Vision The language semantics has been widely exploited in many vision subareas including zero-shot object detection (Bansal et al. 2018), zero-shot recognition (Wang, Ye, and Gupta 2018), few-shot object detection (Zhu et al. 2021) and HOI detection (Peyre et al. 2019; Xu et al. 2019). A typical method that the above works adopt is to push vision space closer to semantic space, which fails to provide positive results when the vision model is strong. In this paper, we present a method that utilizes language semantics with dataset-specific prior knowledge, which is better than the supervision from vanilla multi-modal joint embeddings (Xu et al. 2019).

Cross-Modal Interaction Cross-modal interaction is widely studied in cross-modal retrieval (Wang et al. 2019b; Liu et al. 2019) and VQA (Gao et al. 2019; Jiang et al. 2020). A typical practice of cross-modal interaction is to use the attention mechanism (Vaswani et al. 2017) to achieve cross-modal global context aggregation. While in our framework, the features in one modality will be dominated by *background* class (Carion et al. 2020), which may degrade the quality of cross-modal context aggregation. Hence, we drop this paradigm. Instead, we introduce Cross-Modal Calibration to better incorporate cross-modal supervision for HOI.

Methodology

In this section, we detail Object-guided Cross-modal Calibration Network (OCN) by introducing its overall pipeline, details of Verb Semantic Model (VSM) and Cross-Modal Calibration (CMC), and the training and inference strategy.

Overall Pipeline

The overall pipeline of OCN is shown in Figure 1, which is mainly composed of three parts: a Vision Model (VM), a Verb Semantic Model (VSM) and CMC that benefits from two-modal features. For VM, we refer to the DETR-based HOI models (Zou et al. 2021; Tamura, Ohashi, and Yoshinaga 2021) due to their compact structure. Given an input image, VM adopts a backbone (*e.g.* ResNet-50 (He et al. 2016)) to extract image features. We add fixed positional encoding (Gehring et al. 2017) to the reduced features. Then the features are flattened and fed to a Transformer encoder (Vaswani et al. 2017). We define a set of learned HOI queries $\mathcal{Q} = \{q_i \in \mathbb{R}^D\}_{i=1}^{N_q}$ to perform decoding, producing decoded query features $\bar{\mathcal{Q}} = \{\bar{q}_i \in \mathbb{R}^D\}_{i=1}^{N_q}$, where N_q, D denote the number of queries and the dimension of queries respectively. Note that one HOI query is responsible for the detection of one HOI triplet $\langle human, verbs, object \rangle$. Thus, the decoded feature \bar{q}_i is then fed to independent Feed-Forward Networks (FFN) to predict a human box, an object box and an object class.

VSM mainly consists of a verb semantic reasoning module. The inputs of VSM are word embeddings. To partly

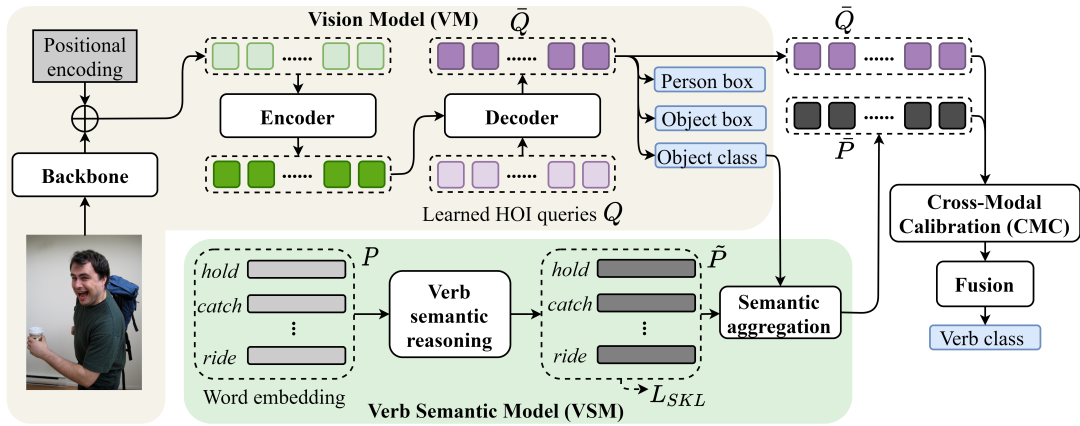


Figure 1: The overall pipeline of OCN. Generally, it consists of VM, VSM and CMC. Input word embeddings are shared across different images for a given dataset.

close the gap between initial word embeddings and the HOI dataset, a newly-proposed SKL loss is used to optimize the embeddings with verb co-occurrence priors. To benefit the verb prediction, we use object-guided verb semantic aggregation to generate a set of semantic embeddings \tilde{P} . Combined with \bar{Q} and \tilde{P} , we can perform CMC. After fusing features from two modalities, another FNN is applied to predict verb classes.

Verb Semantic Model

VSM aims to create a semantic space that is aligned with the training set. Word embeddings from pretrained models *e.g.* GloVe (Pennington, Socher, and Manning 2014) inherently contain co-occurrence priors aligned with certain data *e.g.* Google News, which may not fit in with our dataset. In VSM, we explicitly inject dataset-specific co-occurrence priors via verb semantic reasoning and SKL Loss.

To facilitate the semantic space to better close the discrepancy, we adopt a graph formulation for projection. Suppose the initialized word embeddings are $\mathbf{P} = \{\mathbf{p}_i \in \mathbb{R}^{D_p}\}_{i=1}^{N_p}$, where N_p, D_p denote the number of verbs in the HOI dataset and the dimension of word embeddings. Note that word embeddings are ℓ_2 -normed. We obtain $\tilde{\mathbf{P}} = \{\tilde{\mathbf{p}}_i \in \mathbb{R}^{D_p}\}_{i=1}^{N_p}$ after verb semantic reasoning by

$$r_{ij} = \frac{\theta(\mathbf{p}_i)^T \phi(\mathbf{p}_j)}{\sqrt{D}}; \quad r'_{ij} = \text{softmax}_j(r_{ij}) \quad (1)$$

$$\tilde{\mathbf{p}}_i = \sigma \left(\sum_{j=1}^{N_p} r'_{ij} \mathbf{W}_{p1} \mathbf{p}_j \right) + \mathbf{W}_{p2} \mathbf{p}_i \quad (2)$$

where σ is the non-linear activation function ReLU; θ, ϕ are both linear projections that embed \mathbf{p}_i into D -dimension space; r_{ij} denotes pairwise relation; softmax_j denotes softmax function conducted along index j ; $\mathbf{W}_{p1}, \mathbf{W}_{p2} \in \mathbb{R}^{D \times D_p}$ are linear projections for value embedding and residual connection.

Similarity KL To inject priors into the semantics $\tilde{\mathbf{P}}$, SKL is designed to optimize the adjacency matrix of $\tilde{\mathbf{P}}$ to obey

the co-occurrence distribution for a given HOI dataset. Due to the verb imbalanced problem, the naive joint distribution of verb pairs will be dominated by head classes. Inspired by (You et al. 2020), we generate a symmetrized conditional distribution. We define the verb from the training set as a set $V = \{v_i\}_{i=1}^{N_p}$ and its conditional probability as $\mathbf{C} = \{c_{ij} = \mathbb{P}(v_j|v_i)|i, j = 1, 2, \dots, N_p; i \neq j\}$. The symmetrized conditional distribution $\hat{\mathbf{C}}$ can be obtained by

$$\hat{c}_{ij} = \frac{c_{ij} + c_{ji}}{2N_p} \quad (3)$$

where \hat{c}_{ij} denotes the symmetrized probability for c_{ij} . Note that the denominator $2N_p$ is a normalizing factor to make $\hat{\mathbf{C}}$ sum to one. We use ℓ_2 -norm to normalize $\tilde{\mathbf{p}}_i$. The adjacency matrix $\mathbf{A} = \{a_{ij} = \mathbb{P}(v_j|v_i)|i, j = 1, 2, \dots, N_p; i \neq j\}$ of the semantics $\tilde{\mathbf{P}}$ can be obtained by

$$a_{ij} = \frac{\exp(\tilde{\mathbf{p}}_i^T \tilde{\mathbf{p}}_j / \tau)}{\sum_{k=1}^{N_p} \sum_{l=1, l \neq k}^{N_p} \exp(\tilde{\mathbf{p}}_k^T \tilde{\mathbf{p}}_l / \tau)} \quad (4)$$

where τ is a temperature parameter that scales the softmax distribution of inner products of the normalized semantics (Wu et al. 2018). Our aim is to push close distribution \mathbf{A} to distribution $\hat{\mathbf{C}}$. We utilize the KL-divergence to achieve this purpose, which is formulated as

$$\mathcal{L}_{\text{SKL}} = \mathbb{E}_{\hat{\mathbf{C}}}[\log(\hat{\mathbf{C}}) - \log(\mathbf{A})] \quad (5)$$

By the utilization of \mathcal{L}_{SKL} , the semantics are forced to obey the co-occurrence priors.

Semantic Aggregation To facilitate the verb prediction, we propose object-guided verb semantic aggregation. The basic intuition is that given an object class for one specific HOI triplet, it indicates what the verb might be or might not be based on the object-verb co-occurrence priors. To realize this, we first collect the conditional priors $\mathbf{S} = \{s_{ij} = \mathbb{P}(v_j|o_i)|i = 1, 2, \dots, N_o + 1; j = 1, 2, \dots, N_p\}$ from the training set, where $o_i \in \mathbf{O} = \{o_i\}_{i=1}^{N_o+1}$, N_o denotes the number of object classes and the additional 1 denotes the *background* class. Note that $s_{ij}|_{i=N_o+1}$ is manually set to be a uniform distribution, as the *background* class

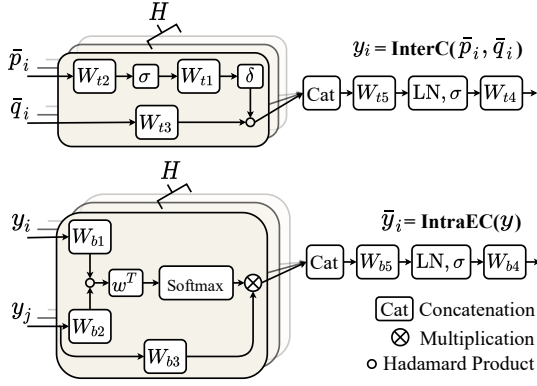


Figure 2: Example visualizations for $y_i = \text{InterC}(\bar{p}_i, \bar{q}_i)$ and $\bar{y}_i = \text{IntraEC}(y)$. softmax is used along index j .

gives no clue about the verb classes. To overcome the effect caused by false prediction of object classes and to alleviate the long-tailed distribution of verbs, we use Laplacian Smoothing (Zhai and Lafferty 2004) to smooth the conditional distribution s_{ij} as

$$\hat{s}_{ij} = \frac{s_{ij} + \beta/N_p}{\sum_{k=1}^{N_p} (s_{ik} + \beta/N_p)} = \frac{s_{ij} + \beta/N_p}{1 + \beta} \quad (6)$$

where \hat{s}_{ij} denotes the smoothed object-verb priors for s_{ij} ; β is a hyper-parameter for smoothing.

Given a decoded HOI query set Q , we predict their object classes by FFN. Combined with object-verb priors, we manage to aggregate a set of $\bar{P} = \{\bar{p}_i \in \mathbb{R}^D\}_{i=1}^{N_q}$, which indicates a verb guess based on the smoothed priors. Note that Q and \bar{P} are one-to-one match, and the cardinalities of \bar{P} and P are different. The verb semantic aggregation process can be formulated as

$$\bar{q}_i^{(o)} = \text{FFN}(\bar{q}_i); \quad \bar{p}_i = \sum_{j=1}^{N_p} \hat{s}_{\bar{q}_i^{(o)} j} \tilde{p}_j \quad (7)$$

where $\bar{q}_i^{(o)}$ denotes the object class of query \bar{q}_i predicted using FFN. $\hat{s}_{\bar{q}_i^{(o)} j}$ is the smoothed $\mathbb{P}(v_j | o_{\bar{q}_i^{(o)}})$. Eq.7 utilizes the object class as a cue, to aggregate verb semantic embeddings based on the smoothed priors $\hat{s}_{\bar{q}_i^{(o)} j}$.

Cross-Modal Calibration

The obtained set Q and P contain visual cues and semantic cues for HOI verbs respectively. To overcome the static semantic embedding problem, we propose to calibrate one modality guided by the other modality. Previous work globally aggregate features from the other modality (Gao et al. 2019) but both modality in our framework is dominated by *background* class, which will not result in a fine cross-modal aggregation. Inspired by (Hu, Shen, and Sun 2018; Vaswani et al. 2017), we propose to perform cross-modal feature calibration in multiple subspaces, which includes InterC and IntraEC illustrated in Figure 2. Below we will exemplify using \bar{p}_i to calibrate \bar{q}_i , denoted as $y_i = \text{InterC}(\bar{p}_i, \bar{q}_i)$ and performing intra-modal reasoning, denoted as $\bar{y}_i = \text{IntraEC}(y)$. The generation of $x_i = \text{InterC}(\bar{q}_i, \bar{p}_i)$ and $\bar{x}_i = \text{IntraEC}(x)$ is similarly conducted.

Inter-Modal Calibration To perform InterC, we project \bar{p}_i into a subspace and then excite other modality’s feature \bar{q}_i in the subspace. We formulate it as

$$e_i = \delta(W_{t1}\sigma(W_{t2}\bar{p}_i)) \circ W_{t3}\bar{q}_i \quad (8)$$

where $W_{t2}, W_{t3} \in \mathbb{R}^{\frac{D}{H} \times D}$ project \bar{p}_i and \bar{q}_i into $\frac{D}{H}$ -dimension subspaces; $W_{t1} \in \mathbb{R}^{\frac{D}{H} \times \frac{D}{H}}$ projects the σ -activated feature into $\frac{D}{H}$ -dimension subspaces; δ is the sigmoid function; \circ denotes Hadamard product. We project the \bar{q}_i in rather low dimension in order to perform multi-head calibration with acceptable computational cost. We can perform the calibration as

$$y_i = \bar{q}_i + W_{t4}\sigma(\text{LN}(W_{t5}\text{Cat}(e_i^{(1)}, \dots, e_i^{(H)}))) \quad (9)$$

where H denotes the number of independent InterC heads; superscript is added to e_i , denoted as different heads; LN denotes LayerNorm (Ba, Kiros, and Hinton 2016); Cat denotes concatenation; $W_{t4}, W_{t5} \in \mathbb{R}^{D \times D}$.

Intra-Modal Enhanced Calibration Since InterC calibrates the features, it is intuitive to incorporate a global reasoning module IntraEC to restore the global context. In practice, we adopt bilinear pooling (Kim et al. 2016; Yuan and Ni 2021) to infer intra-modal relation as

$$f_{ij} = w^T(W_{b1}y_i) \circ (W_{b2}y_j) \quad (10)$$

$$f'_{ij} = \text{softmax}_j(f_{ij}); \quad g_{ij} = f'_{ij}W_{b3}y_j \quad (11)$$

where $w \in \mathbb{R}^D$, $W_{b1}, W_{b2} \in \mathbb{R}^{D \times D}$ and $W_{b3} \in \mathbb{R}^{\frac{D}{H} \times D}$. Also, we extend IntraEC into a multi-head manner for better representation ability:

$$\bar{y}_i = y_i + W_{b4}\sigma(\text{LN}(W_{b5} \sum_{j=1}^{N_q} \text{Cat}(g_{ij}^{(1)}, \dots, g_{ij}^{(H)}))) \quad (12)$$

where $W_{b4}, W_{b5} \in \mathbb{R}^{D \times D}$; H denotes the number of independent IntraEC heads, which is identical to InterC.

Modality Fusion After CMC is conducted, we obtain two sets of verified features from two modalities. We adopt the fusion strategy (Zhang, Hare, and Prügél-Bennett 2018) to fuse \bar{x}_i and \bar{y}_i as

$$z_i = \sigma(W_x\bar{x}_i + W_y\bar{y}_i) - (W_x\bar{x}_i - W_y\bar{y}_i)^2 \quad (13)$$

where $W_x, W_y \in \mathbb{R}^{D \times D}$. The obtained feature set $Z = \{z_i\}_{i=1}^{N_q}$ serves as the feature for predicting verbs by FFN.

Training and Inference Strategy

Similar to all DETR-based methods, we formulate HOI detection as a set prediction problem (Carion et al. 2020). To train the end-to-end model, we need to **i)** match ground truth (GT) HOI triplets with the predicted HOI triplets, **ii)** calculate losses for VM and VSM. The matching process is conducted with Hungarian algorithm (Kuhn 1955). If we denote GT HOI triplet set as $G = \{g_i\}_{i=1}^{N_q}$ (padded with no HOI triplets in order to match) and the predicted HOI set as $M = \{m_i\}_{i=1}^{N_q}$, the bipartite matching can be formulated as

$$\hat{\psi} = \arg \min_{\psi \in \Psi_{N_q}} \sum_{i=1}^{N_q} \mathcal{H}_{\text{cost}}(g_i, m_{\psi(i)}) \quad (14)$$

\mathcal{L}_{SKL}	VSM	InterC	IntraEC	Full	Rare	Non-Rare
	Base VM			29.15	22.20	31.23
✓	✓	✓	✓	30.91	25.56	32.51
	✓	✓	✓	30.54	24.33	32.39
✓	✓			29.88	23.50	31.79
✓	✓	✓		30.51	25.15	32.12
✓	✓		✓	30.40	24.71	32.10

(a) Ablation study of proposed modules and loss.

#Heads H	Full	Rare	Non-Rare
1	30.58	25.65	32.05
2	30.91	25.56	32.51
4	30.82	24.45	32.73
8	30.65	24.93	32.36

(b) Effect of head numbers for both InterC and IntraEC.

τ	Full	Rare	Non-Rare
0.025	30.70	25.03	32.39
0.050	30.91	25.56	32.51
0.100	30.54	24.67	32.29
0.200	30.72	25.59	32.25

(c) Effect of varying choices for temperature τ in \mathcal{L}_{SKL} .

Table 1: Ablation study on HICO-DET.

where Ψ_{N_q} is the solution space for bipartite matching. The matching cost $\mathcal{H}_{\text{cost}}$ follows (Tamura, Ohashi, and Yoshinaga 2021) and is detailed in the Supplementary Material. The loss to train the model can be denoted as

$$\mathcal{L} = \lambda_1 \mathcal{L}_{SKL} + \lambda_2 \mathcal{L}_{box} + \lambda_3 \mathcal{L}_{GIoU} + \lambda_4 \mathcal{L}_o + \lambda_5 \mathcal{L}_v \quad (15)$$

where \mathcal{L}_{box} denotes ℓ_1 loss for box regression; \mathcal{L}_{GIoU} denotes GIoU loss (Rezatofighi et al. 2019); \mathcal{L}_o denotes Cross-Entropy loss for object class; \mathcal{L}_v denotes loss for verbs. We mainly study Binary Cross-Entropy (BCE) and Focal loss (Lin et al. 2017) as \mathcal{L}_v . λ balances these losses by setting different weights.

During inference, the object class and the bounding boxes (bbox) of the human and the object for one HOI triplet is simply generated from \bar{q}_i . The object score is the maximum object confidence score. The verb score is the multiplication of the object score and the verb score predicted by z_i . A binary mask is used by default to filter out object-verb pairs that are not in the training set.

Experiments

Datasets and Metrics

In this paper, we use two widely-adopted datasets dubbed HICO-DET (Chao et al. 2015) and V-COCO (Gupta and Malik 2015). HICO-DET contains 37,536 training images and 9,515 testing images, in which 600 $\langle \text{verb}, \text{object} \rangle$ unique interaction types are defined out of 117 verb classes and 80 object classes. We evaluate on the test set by interaction mAP (%) over three sets: **i**) Full set (all 600 interactions), **ii**) Rare set (138 interactions with less than 10 training samples), **iii**) Non-Rare set (462 interactions with 10 or more training samples). We evaluate our model under *Default* setting for the 3 sets. One HOI triplet is rightly localized when the predicted bboxes of the human and object have Intersection-over-Union (IoU) greater than 0.5 with GT bboxes.

V-COCO contains 2,533 training images, 2,867 validating images and 4,946 testing images, in which interactions are defined upon 25 interactions and 80 object classes. We evaluate on the test set by verb mAP (%) under two scenarios following (Kim et al. 2021; Tamura, Ohashi, and Yoshinaga 2021): **i**) In Scenario1, we need to report when there is no object in the GT HOI triplet, denoted as $AP_{role}^{\#1}$; **ii**) In Scenario2, we can ignore the prediction of the object bbox when the GT HOI triplet is without object, denoted as $AP_{role}^{\#2}$.

Implementation Details

For VM, we use Transformer with a 6-layer encoder and a 6-layer decoder following (Carion et al. 2020). We use parameters of DETR trained on COCO (Lin et al. 2014) as VM’s initialization. We adopt AdamW (Loshchilov and Hutter 2018) to optimize OCN for 80 epochs with a weight decay of 10^{-4} . The learning rate (lr) of the backbone is fixed to 10^{-5} . The lr of other parts starts from 10^{-4} and decays to 10^{-5} after the 60th epoch. We use basic data augmentation to train a robust model, including random crop, random horizontal flipping, image scale and color augmentation following (Carion et al. 2020; Liao et al. 2020). During evaluation, the most confident $K = 100$ HOI triplets are selected to compute mAP.

By default, we use following parameters if not otherwise stated. We set number of queries $N_q = 100$ and the dimension of queries $D = 256$. The head number H for InterC and IntraEC is set to 2. The temperature τ of \mathcal{L}_{SKL} is set to 0.05. The smoothing parameter β is set to 0.1. N_p is the number of verb classes. The verb loss is Focal loss and the backbone is ResNet-50 by default. We use 300-dimension GloVe word embedding (Pennington, Socher, and Manning 2014) as P . The hyper-parameters $\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5$ in \mathcal{L} are set to 1, 2.5, 1, 1, 1 respectively.

Model Analysis

Ablation Study We conduct model analysis on HICO-DET. Note that because VM’s structure **does not change**, the ability to localize HOI triplets stays nearly unchanged. The boost purely comes from better verb-inferring ability.

Module and Loss Ablation To demonstrate the efficacy of our proposed module, we first conduct ablation experiments as shown in Table 1a. By applying our full model, the Rare set can gain 3.36% and the Non-Rare set 1.28%, producing a more balanced result. If removing \mathcal{L}_{SKL} , the model trusts the implicit co-occurrence containing in the initialized word embedding, which is not aligned with our dataset, thus causing performance decline. By applying the VSM together with \mathcal{L}_{SKL} , the performance can be boosted by basic semantic aggregation, indicating that the object-verb prior based semantic aggregation helps to better rank the verbs. By using InterC, two modal features can calibrate features based on the other modality, bringing in performance boost. By using IntraEC, intra-modal features, especially semantic aggregated features \bar{p}_i , can have better global-dependent representations (Vaswani et al. 2017). By

Smoothing β	Full	Rare	Non-Rare
0	30.50	25.00	32.15
0.1	30.91	25.56	32.51
1	30.85	25.43	32.47
10	30.81	25.40	32.43
100	30.36	24.86	32.01
∞	30.18	23.33	32.22

Table 2: Effect of varying β in Eq.6. $\beta = \infty$ denotes $\hat{s}_{ij} = 1/N_p$, without any prior knowledge.

Verb Loss	Model	Full	Rare	Non-Rare
BCE	VM	20.13	13.16	22.22
	VM + MMJE	20.60 _{+0.47}	13.66 _{+0.50}	22.67 _{+0.45}
	OCN	26.42 _{+6.29}	20.79 _{+7.63}	28.11 _{+5.83}
Focal	VM	29.15	22.20	31.23
	VM + MMJE	28.83 _{-0.32}	22.09 _{-0.11}	30.84 _{-0.39}
	OCN	30.91 _{+1.76}	25.56 _{+3.36}	32.51 _{+1.28}

Table 3: Performance analysis on HICO-DET with different verb losses and semantic models.

appending IntraEC to InterC, features in identical modality can restore the global context after calibration, bringing additive performance boost to InterC.

Sensitivity Analysis of Head Numbers InterC and IntraEC are designed to be multi-head. We vary the head numbers for both modules, the results of which are shown in Table 1b. The best choice for the head number is 2.

Sensitivity Analysis of the Temperature Parameter For the optimization of VSM, we vary the temperature τ in \mathcal{L}_{SKL} which is shown in Table 1c. A low value of τ will sharpen the distribution of \hat{A} and thus ease the optimization of \mathcal{L}_{SKL} . The table indicates the optimal value is 0.05.

Sensitivity Analysis of the Smoothed Distribution An appropriate probability smoothing hyper-parameter β in Eq.6 can balance the HOI detection result by smoothed semantic aggregation. We try varying β and results are in Table 2. The table indicates that, **i)** compared to non-smoothed semantic aggregation, properly smoothing the object-verb distribution can slightly boost the performance; **ii)** the smoothing operation is not sensitive to the choice of β , with β ranging from 0.1 to 10 producing similar results. Hence we set $\beta = 0.1$ by default; **iii)** by setting $\beta = \infty$, the model loses the verb prediction orientation, which damages the performance and proves our claim.

OCN Helps More for Poor Verb Predictor Since our model relies on object-verb and verb-verb priors, it injects verb prediction guessing information given an object. We have reasonable speculation that adding proposed modules on a model with a poor verb predictor will have more significant improvements due to this object-guided structure. We conduct an experiment with a loss that has more trouble in inferring verbs (BCE), shown in Table 3. BCE suffers from the problem of imbalanced positive-negative samples. OCN with BCE loss improves upon VM by 6.29%, suggesting that OCN greatly ameliorates the imbalance problem. Even with Focal loss, our model can also make its contribution. More-

Model Mask	Full	Rare	Non-Rare	mR@100
VM ✓	28.89	21.76	31.02	57.83
	29.15 _{+0.26}	22.20 _{+0.44}	31.23 _{+0.21}	65.10 _{+7.27}
OCN ✓	30.82 _{+1.93}	25.48 _{+3.72}	32.42 _{+1.40}	63.86 _{+6.03}
	30.91 _{+2.02}	25.56 _{+3.80}	32.51 _{+1.49}	67.64 _{+9.81}

Table 4: Performance analysis on HICO-DET with different methods of utilizing prior knowledge. "Mask" denotes binary mask used to filter out impossible object-verb pairs.

Method	Backbone	Time	#Params	Full
PPDM (Liao et al. 2020)	HOG104	56ms	194.9M	21.94
HOTR (Kim et al. 2021)	R50	61ms	51.2M	25.10
ASNet (Chen et al. 2021)	R50	56ms	52.5M	28.87
QPIC	R50	38ms	41.5M	29.07
OCN	R50	43ms	43.4M	30.91

Table 5: Computational cost analysis on HICO-DET with Tesla V100. HOG and R are short for Hourglass and ResNet.

over, the Rare set benefits more than the Non-Rare set thanks to our object-guided structure.

Superiority over Multi-Modal Joint Embeddings A common practice of utilizing word embeddings is Multi-Modal Joint Embedding (MMJE) (Xu et al. 2019), which maximizes the similarity between positive vision-semantic pairs and keep negative pairs to a predefined margin. We reimplement MMJE by VSM and \mathcal{L}_{sim} in (Xu et al. 2019) and add MMJE to VM, results of which are shown in Table 3. Comparing VM and VM+MMJE, MMJE will cause negative effects when VM is already strong because it trusts the underlying verb relations in word embeddings which have a discrepancy towards the HOI dataset. However, our method with priors will constantly bring in positive effects.

Superiority over Binary Mask The binary mask can be treated as a stiff way to utilize relation prior knowledge by applying a hard mask to the predicted object-verb pair. Our method can be considered as a more fine-grained method which considers object-verb priors and verb co-occurrence priors. The results are shown in Table 4. In this table, we also provide results with a metric: mean Recall@K (%) (Tang et al. 2020) (mR@K), which averages the recall of HOI interactions with top K evaluation. K is set to 100 in this paper. By applying a stiff method utilizing object-verb prior, although mR@100 goes up by 7.27%, the result of the Full set can only be boosted by 0.26%. While OCN w/o binary mask can boost the Full set by 1.93% and the Rare set by 3.72%, with mR@100 merely going up by 6.03%. It indicates that OCN improves the ranking of HOI interactions greatly while the binary mask struggles to do so, demonstrating the significance of mining the object-guided verb prediction structure. If adding binary masks to OCN, the mAP improvement will be very trivial but with an obvious mR@100 boost, which reemphasizes the binary masks' poor ability.

Computational Cost Analysis We compare the computational cost with competitive one-stage methods, shown in Table 5. It can be seen from the table that our method does

Method	Backbone	Full	Rare	Non-Rare	AP _{role} ^{#1}	AP _{role} ^{#2}
Two-stage						
InteractNet	R50-FPN	9.94	7.16	10.77	40.0	-
GPNN	R152-DCN	13.11	9.34	14.23	44.0	-
iCAN	R50	14.84	10.45	16.15	45.3	52.4
No-Frills*	R152	17.18	12.17	18.68	-	-
PMFNet	R50-FPN	17.46	15.65	18.00	52.0	-
DRG*	R50-FPN	19.26	17.74	19.71	51.0	-
IPNet	HOG104	19.56	12.79	21.58	51.0	-
FCMNet*	R50	20.41	17.34	21.56	53.1	-
PD-Net*	R152-FPN	20.81	15.90	22.28	52.6	-
IDN	R50	23.36	22.47	23.63	53.3	60.3
One-stage						
UnionDet	R50-FPN	17.58	11.72	19.33	47.5	56.2
PPDM	HOG104	21.94	13.97	24.32	-	-
GGNet	HOG104	23.47	16.48	25.60	54.7	-
DIRV	ED-d3	21.78	16.38	23.39	56.1	-
HOTR	R50	25.10	17.34	27.42	55.2	64.4
HOITransformer	R50	23.46	16.91	25.41	52.9	-
QPIC [†]	R50	29.07	21.85	31.23	61.8	64.1
ASNet	R50	28.87	24.25	30.25	53.9	-
Ours-OCN*	R50	30.91	25.56	32.51	64.2	66.3
	R101	31.43	25.80	33.11	65.3	67.1

Table 6: Comparisons with state-of-the-arts on HICO-DET and V-COCO. Full, Rare and Non-Rare columns are reported on HICO-DET and AP_{role}^{#1}, AP_{role}^{#2} on V-COCO. R, HOG and ED are short for ResNet, Hourglass and Efficient-Det respectively. * denotes utilization of word embeddings. † denotes reproduction using COCO pre-trained parameters.

well in speed-performance trade-off. To be more specific, VM takes 38ms to run that is identical to QPIC. VSM takes 0.6ms to run. InterC and IntraEC take 4.3ms to run. Note that during inference, we can slightly speed it up by inferring VSM once and storing \tilde{P} in the memory for a given dataset. Our model achieves the best performance with acceptable cost adding to VM.

Object-Conditioned Verb Distribution Analysis To analyze how the object-guided hierarchy helps the verb prediction, we conduct a case study with different methods’ verb distribution on the rightly localized HOI triplets whose objects are predicted as *toaster*, the figure of which is illustrated in Figure 3. Note that VM and OCN have the same vision structure, thus having identical detection abilities. VM predicts the verb more uniformly. OCN, equipped with object-guided priors, predicts the verb distribution aligned more with the distribution of the training set. We also measure the mean Pearson Correlation Coefficient (mPCC) (Benesty et al. 2009) of the predicted object-conditioned verb distribution to the training object-conditioned verb distribution (“mean” averages different objects). It turns out that OCN has a higher mPCC (0.636) than VM does (0.476), indicating the object-guided priors help the verb prediction.

t-SNE Visualization of the Semantic Space To observe SKL loss’s role, we use t-SNE (Maaten and Hinton 2008) to visualize the input GloVe embedding set P and the optimized set \tilde{P} , illustrated in Figure 4. In Figure 4, we high-

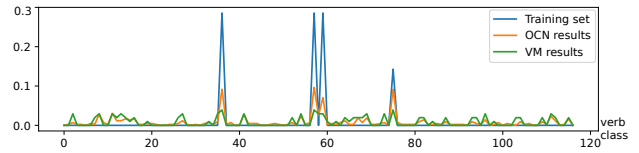


Figure 3: A case visualization of *toaster*-conditioned verb distribution with different methods. Better view in color.

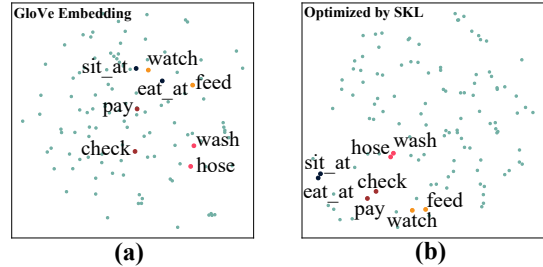


Figure 4: Visualization of the semantic space on HICO-DET.

light four word pairs with different colors that highly co-occur in HICO-DET. The original GloVe embedding spreads in the space while the optimization of SKL loss pushes semantically similar verbs closer, which decreases the discrepancy between the implicit priors in GloVe and the verb co-occurrence prior in HICO-DET.

Comparisons with State-of-the-Arts

We compare our results with previous state-of-the-arts on HICO-DET in Table 6. The table indicates that **i)** All previous models trail our model by a considerable margin. None of the one-stage models possesses the object-guided structure, which shows its superiority. **ii)** Comparing with previous methods that utilize word embeddings, our method greatly surpasses them by a considerable margin and maintains an end-to-end manner. **iii)** Comparing to other DETR-based models *e.g.* HOTR, HOITransformer and QPIC, our method is a more balanced detector thanks to the smoothed prior knowledge.

We compare our results with previous state-of-the-arts on V-COCO in Table 6. The table indicates our method’s more superior performances compared to methods w/ or w/o word embeddings. Our method surpasses previous best AP_{role}^{#1} by 2.4% mAP and best AP_{role}^{#2} by 1.9%.

Conclusions and Future Work

In this paper, we propose to facilitate end-to-end HOI detection models with object-guided priors. In practice, we resort to a more aligned semantic space and propose to perform cross-modal calibration. Similar thoughts can also facilitate other visual relation detection problems like scene graph generation and video HOI. However, our method can run into trouble when extending to scenarios like zero-shot HOI detection, which is left for future exploration.

Acknowledgements

We would like to appreciate Rong Jin for carefully revising the manuscript, Mingqian Tang and Jianwen Jiang for discussing the ideas, and anonymous reviewers for their valuable feedback. This work was supported by National Natural Science Foundation of China Grant No. 62173298.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bansal, A.; Rambhatla, S. S.; Shrivastava, A.; and Chellappa, R. 2020. Detecting human-object interactions via functional generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10460–10469.
- Bansal, A.; Sikka, K.; Sharma, G.; Chellappa, R.; and Divakaran, A. 2018. Zero-shot object detection. In *European Conference on Computer Vision*, 384–400.
- Benesty, J.; Chen, J.; Huang, Y.; and Cohen, I. 2009. *Noise reduction in speech processing*, volume 2. Springer Science & Business Media.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers. *arXiv preprint arXiv:2005.12872*.
- Chao, Y.-W.; Liu, Y.; Liu, X.; Zeng, H.; and Deng, J. 2018. Learning to detect human-object interactions. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 381–389. IEEE.
- Chao, Y.-W.; Wang, Z.; He, Y.; Wang, J.; and Deng, J. 2015. Hico: A benchmark for recognizing human-object interactions in images. In *Proceedings of the IEEE International Conference on Computer Vision*, 1017–1025.
- Chen, M.; Liao, Y.; Liu, S.; Chen, Z.; Wang, F.; and Qian, C. 2021. Reformulating hoi detection as adaptive set prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9004–9013.
- Chen, T.; Yu, W.; Chen, R.; and Lin, L. 2019. Knowledge-embedded routing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6163–6171.
- Gao, C.; Xu, J.; Zou, Y.; and Huang, J.-B. 2020. Drg: Dual relation graph for human-object interaction detection. In *European Conference on Computer Vision*, 696–712. Springer.
- Gao, C.; Zou, Y.; and Huang, J.-B. 2018. ican: Instance-centric attention network for human-object interaction detection. *arXiv preprint arXiv:1808.10437*.
- Gao, P.; Jiang, Z.; You, H.; Lu, P.; Hoi, S. C. H.; Wang, X.; and Li, H. 2019. Dynamic Fusion With Intra- and Inter-Modality Attention Flow for Visual Question Answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional sequence to sequence learning. In *International Conference on Machine Learning*, 1243–1252. PMLR.
- Gkioxari, G.; Girshick, R.; Dollár, P.; and He, K. 2018. Detecting and recognizing human-object interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8359–8367.
- Gupta, S.; and Malik, J. 2015. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*.
- Gupta, T.; Schwing, A.; and Hoiem, D. 2019. No-frills human-object interaction detection: Factorization, layout encodings, and training techniques. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9677–9685.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7132–7141.
- Jiang, J.; Chen, Z.; Lin, H.; Zhao, X.; and Gao, Y. 2020. Divide and conquer: Question-guided spatio-temporal contextual attention for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11101–11108.
- Johnson, J.; Krishna, R.; Stark, M.; Li, L.-J.; Shamma, D.; Bernstein, M.; and Fei-Fei, L. 2015. Image retrieval using scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3668–3678.
- Kim, B.; Choi, T.; Kang, J.; and Kim, H. J. 2020a. Uniondet: Union-level detector towards real-time human-object interaction detection. In *European Conference on Computer Vision*, 498–514. Springer.
- Kim, B.; Lee, J.; Kang, J.; Kim, E.-S.; and Kim, H. J. 2021. HOTR: End-to-End Human-Object Interaction Detection with Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 74–83.
- Kim, D.-J.; Sun, X.; Choi, J.; Lin, S.; and Kweon, I. S. 2020b. Detecting human-object interactions with action co-occurrence priors. In *European Conference on Computer Vision*, 718–736. Springer.
- Kim, J.-H.; On, K.-W.; Lim, W.; Kim, J.; Ha, J.-W.; and Zhang, B.-T. 2016. Hadamard product for low-rank bilinear pooling. *arXiv preprint arXiv:1610.04325*.
- Kuhn, H. W. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2): 83–97.
- Li, Y.-L.; Liu, X.; Wu, X.; Li, Y.; and Lu, C. 2020. HOI Analysis: Integrating and Decomposing Human-Object Interaction. *Advances in Neural Information Processing Systems*, 33.
- Li, Y.-L.; Zhou, S.; Huang, X.; Xu, L.; Ma, Z.; Fang, H.-S.; Wang, Y.; and Lu, C. 2019. Transferable interactiveness knowledge for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3585–3594.
- Liao, Y.; Liu, S.; Wang, F.; Chen, Y.; Qian, C.; and Feng, J. 2020. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 482–490.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 740–755. Springer.
- Lin, X.; Ding, C.; Zeng, J.; and Tao, D. 2020. GPS-Net: Graph Property Sensing Network for Scene Graph Generation. *arXiv preprint arXiv:2003.12962*.
- Liu, X.; Wang, Z.; Shao, J.; Wang, X.; and Li, H. 2019. Improving referring expression grounding with cross-modal attention-guided erasing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1950–1959.

- Liu, Y.; Chen, Q.; and Zisserman, A. 2020. Amplifying key cues for human-object-interaction detection. In *European Conference on Computer Vision*, 248–265. Springer.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov): 2579–2605.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 3111–3119.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Peyre, J.; Laptev, I.; Schmid, C.; and Sivic, J. 2019. Detecting unseen visual relations using analogies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1981–1990.
- Qi, S.; Wang, W.; Jia, B.; Shen, J.; and Zhu, S.-C. 2018. Learning Human-Object Interactions by Graph Parsing Neural Networks. In *European Conference on Computer Vision*.
- Rahman, S.; Khan, S.; Barnes, N.; and Khan, F. S. 2020. Any-shot object detection. In *Proceedings of the Asian Conference on Computer Vision*.
- Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 658–666.
- Tamura, M.; Ohashi, H.; and Yoshinaga, T. 2021. 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*, 10410–10419.
- Tang, K.; Niu, Y.; Huang, J.; Shi, J.; and Zhang, H. 2020. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3716–3725.
- Ulutun, O.; Iftekhar, A.; and Manjunath, B. S. 2020. 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*, 13617–13626.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wang, H.; Zheng, W.-s.; and Yingbiao, L. 2020. Contextual heterogeneous graph network for human-object interaction detection. In *European Conference on Computer Vision*, 248–264. Springer.
- Wang, T.; Anwer, R. M.; Khan, M. H.; Khan, F. S.; Pang, Y.; Shao, L.; and Laaksonen, J. 2019a. Deep contextual attention for human-object interaction detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5694–5702.
- Wang, X.; Ye, Y.; and Gupta, A. 2018. Zero-shot recognition via semantic embeddings and knowledge graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6857–6866.
- Wang, Z.; Liu, X.; Li, H.; Sheng, L.; Yan, J.; Wang, X.; and Shao, J. 2019b. Camp: Cross-modal adaptive message passing for text-image retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5764–5773.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3733–3742.
- Xu, B.; Wong, Y.; Li, J.; Zhao, Q.; and Kankanhalli, M. S. 2019. Learning to detect human-object interactions with knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yao, T.; Pan, Y.; Li, Y.; and Mei, T. 2018. Exploring visual relationship for image captioning. In *European Conference on Computer Vision*, 684–699.
- You, R.; Guo, Z.; Cui, L.; Long, X.; Bao, Y.; and Wen, S. 2020. Cross-modality attention with semantic graph embedding for multi-label classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12709–12716.
- Yuan, H.; and Ni, D. 2021. Learning Visual Context for Group Activity Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 3261–3269.
- Yuan, H.; Ni, D.; and Wang, M. 2021. Spatio-Temporal Dynamic Inference Network for Group Activity Recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7476–7485.
- Zhai, C.; and Lafferty, J. 2004. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems (TOIS)*, 22(2): 179–214.
- Zhang, Y.; Hare, J.; and Prügel-Bennett, A. 2018. Learning to count objects in natural images for visual question answering. *arXiv preprint arXiv:1802.05766*.
- Zhong, X.; Ding, C.; Qu, X.; and Tao, D. 2020. Polysemy deciphering network for human-object interaction detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, 69–85. Springer.
- Zhong, X.; Qu, X.; Ding, C.; and Tao, D. 2021. Glance and Gaze: Inferring Action-aware Points for One-Stage Human-Object Interaction Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13234–13243.
- Zhou, P.; and Chi, M. 2019. Relation Parsing Neural Network for Human-Object Interaction Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Zhu, C.; Chen, F.; Ahmed, U.; Shen, Z.; and Savvides, M. 2021. Semantic relation reasoning for shot-stable few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8782–8791.
- Zou, C.; Wang, B.; Hu, Y.; Liu, J.; Wu, Q.; Zhao, Y.; Li, B.; Zhang, C.; Zhang, C.; Wei, Y.; et al. 2021. End-to-end human object interaction detection with hoi transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11825–11834.