

Exploring the Potential of Large Foundation Models for Open-Vocabulary HOI Detection

Ting Lei Shaofeng Yin Yang Liu*

Wangxuan Institute of Computer Technology, Peking University
{ting_lei, yangliu}@pku.edu.cn yin.shaofeng@stu.pku.edu.cn


Abstract

Open-vocabulary human-object interaction (HOI) detection, which is concerned with the problem of detecting novel HOIs guided by natural language, is crucial for understanding human-centric scenes. However, prior zero-shot HOI detectors often employ the same levels of feature maps to model HOIs with varying distances, leading to suboptimal performance in scenes containing human-object pairs with a wide range of distances. In addition, these detectors primarily rely on category names and overlook the rich contextual information that language can provide, which is essential for capturing open vocabulary concepts that are typically rare and not well-represented by category names alone. In this paper, we introduce a novel end-to-end open vocabulary HOI detection framework with conditional multi-level decoding and fine-grained semantic enhancement (CMD-SE), harnessing the potential of Visual-Language Models (VLMs). Specifically, we propose to model human-object pairs with different distances with different levels of feature maps by incorporating a soft constraint during the bipartite matching process. Furthermore, by leveraging large language models (LLMs) such as GPT models, we exploit their extensive world knowledge to generate descriptions of human body part states for various interactions. Then we integrate the generalizable and fine-grained semantics of human body parts to improve interaction recognition. Experimental results on two datasets, SWIG-HOI and HICO-DET, demonstrate that our proposed method achieves state-of-the-art results in open vocabulary HOI detection. The code and models are available at <https://github.com/l1ttt/pku/CMD-SE-release>.


1. Introduction

Human-object interaction (HOI) detection aims to localize the interacting human and object pairs and then recognize

*Corresponding author



(a) Performance comparison between our method and THID on HOIs with different distances on the open-vocabulary SWIG-HOI dataset.



(b) t-SNE visualization of HOI concepts in semantic space.

Figure 1. (a) Previous method (THID) suffers from severe performance drop on HOIs with larger distances in the open-vocabulary setting. (b) Compared with HOI categorical names, body parts' descriptions could better recognize the correlation of human postures between different actions. For instance, the action of hurling and picking typically involves extended arms, whereas kicking is characterized by extended legs.

their interactions, *i.e.* a set of <human, object, action> triplets. Most existing HOI detection algorithms [4, 8, 24, 32, 35, 58, 61, 63–65] are unscalable in terms of vocabulary size, *i.e.*, they treat interactions as discrete labels and train a classifier based on a predefined category space. However, due to the combinatorial nature of interactions, it is impractical to create a data collection to include all possible HOIs, especially when the action and object category space becomes large. This motivates us to study a transferable HOI detector that can detect any HOI given its class name.

Recent works [1, 16–18, 21] propose to utilize compositional learning to enhance the generalization ability of HOI detectors, particularly for unseen interactions. Their core idea is to decompose the interaction into action and object, followed by data augmentation to create novel combinations. However, lacking the help of semantics, the above methods can only detect a small set of HOIs predefined in the dataset. Besides, the aforementioned methods may struggle to detect rare concepts, which are more likely to represent new Open Vocabulary (OV) concepts not included in the annotated data during training. Other works [34, 39, 56, 57] propose to incorporate language priors in zero-shot HOI detection, transforming one-hot HOI labels into natural language supervision by CLIP [45].

Despite the progress, previous methods still face two major challenges in solving the open-vocabulary HOI detection task. First, the distances between interactive human-object pairs often exhibit diverse spatial distances. As a result, modeling HOIs using uniform levels of feature maps may lead to suboptimal performance [25]. As shown in Figure 1a, the existing state-of-the-art approach THID suffers from a severe performance drop on HOIs with large distances. Second, previous detectors rely on category names and overlook the rich contextual information that language can provide. This oversight is particularly pertinent in the context of open vocabulary, where rare concepts are prevalent. In such cases, text embeddings for rare concepts may not be reliable and may fail to reveal the visual appearance similarities between them. Consequently, these detectors may face challenges in capturing the intrinsic connections among different HOI concepts and in generalizing this knowledge from common categories to rare or unseen categories. For example, considering three actions: hurling, picking, and kicking, we can visually discern that hurling and picking involve extended arms and open hands, while kicking is characterized by extended legs and feet. However, in the semantic space depicted on the left side of Figure 1b, the distances between the categorical names of these three actions’ HOIs are equivalent, undermining the model’s capacity to comprehend interactions in an open-vocabulary scenario.

According to the challenges mentioned above, we build an end-to-end open vocabulary HOI detector with large foundation models, which not only models human-object pairs with different distances by utilizing different levels of feature maps, but also exploits the generalizable and fine-grained semantics of human body parts to distinguish various interaction concepts. *Firstly*, to deal with the problem of suboptimal performance with single-level feature maps, we propose the utilization of multi-level feature maps tailored to HOIs at different distances. Our strategy involves decoding HOIs in parallel from these multi-level feature maps. Specifically, we guide the low-level fea-

ture maps to correspond with human-object pairs that have a smaller H-O distance, and the opposite for higher-level maps. This targeted matching is facilitated by introducing an additional constraint during the bipartite matching process [27]. *Secondly*, we propose to utilize the generalizable and fine-grained semantics of human body parts besides category names for interaction recognition. By focusing on body parts, which are more fundamental and elementary compared to the human body as a whole, we aim to enable a more generalizable visual perception across seen and unseen scenarios. Inspired by the remarkable world knowledge on a variety of topics of large language models (LLMs) like GPT models [2, 43, 44], we utilize it as an implicit knowledge base to exploit the generalizable and fine-grained semantics of human body parts to enhance the understanding of interaction concepts. Specifically, we query a LLM with natural language to let it describe the states of human body parts given a HOI. Then, we utilize the generated descriptions of human body parts for interaction recognition, guiding the model in a fine-grained semantic space. As illustrated on the right side of Figure 1b, body parts’ descriptions can align the two closely related actions (hurling and picking), which share common patterns in body parts involvement. Simultaneously, this enhancement effectively distinguishes them from kicking, thereby amplifying the model’s ability to comprehend a broader range of interactions.

Our contributions are summarized below. (1) We propose to utilize different levels of feature maps of VLM to model HOIs with varying distances in the open-vocabulary scenario through a conditional matching mechanism. (2) We incorporate the generalizable and fine-grained semantics of human body parts obtained by querying LLM to enhance the understanding of a large vocabulary of interactions. (3) Experiments on two datasets, SWIG-HOI and HICO-DET validate that the proposed method can achieve state-of-the-art results on Open-Vocabulary HOI detection.

2. Related Work

2.1. Generic HOI Detection

According to the network architecture design, previous HOI detection methods can be categorized into two-stage [4, 11, 12, 28, 30, 41, 63, 64] and one-stage [7, 13, 23, 25, 26, 33, 51, 61, 68] paradigms. Two-stage methods usually apply an object detector first to detect humans and objects, followed by specifically designed modules for human-object association and interaction recognition. They typically use multi-stream [14, 18, 31, 35] or graph-based [52, 59] methods to support interaction understanding. In contrast, one-stage methods typically employ multitask learning to perform instance detection and interactive recognition jointly [24, 33, 50, 61]. Despite the progress, stan-

standard HOI detection treats interactions as discrete labels and learns a classifier based on a predefined category space, lacking the ability to detect numerous potential unseen interactions.

2.2. Vision-Language Modeling in HOI Detection

Although previous HOI detectors have achieved moderate success, they often treat interactions as discrete labels and ignore the richer semantic text information in triplet labels. More recently, some researchers [19, 25, 34, 39, 56, 57, 67, 70] have investigated the generalizable HOI detector with Vision-Language Modeling. Among them, PD-Net [67] and SSRT [19] propose to aggregate language prior features into the interaction recognition. RLIP [70] proposes a pre-training strategy for HOI detection based on image captions. GEN-VLKT [34] and HOICLIP [39] employ the CLIP visual encoder to guide the learning of interaction representation and utilize CLIP text embeddings of prompted HOI labels to initialize the classifier. THID [56] proposes a HOI sequence parser to detect multiple interactions and first gets promising results on the recent open-vocabulary HOI detection dataset [55]. OpenCat [66] leverages massive amounts of weakly supervised data and proposes several proxy tasks for HOI pre-training based on CLIP. However, previous methods suffer from the limitation of utilizing the same levels of feature maps to model HOIs with varying distances, resulting in suboptimal performance. Additionally, these methods overlook the valuable and generalizable semantics of human body parts, which can enhance the understanding of interactions at a fine-grained level.

2.3. Leverage LLM for Text Classifier

Distinct from traditional methods like manually crafting descriptions [15, 46] or utilizing external databases such as Wikipedia [9, 37] or the WordNet hierarchy [10, 47, 49], recent studies have demonstrated the effectiveness of Large Language Models (LLMs) in generating descriptive prompts for the classification and detection tasks. Some works [36, 38, 42] utilize GPT-3 [2] to create detailed sentences that capture visual concepts for recognizing specific categories. Others [40, 48] propose using LLMs to generate semantic hierarchies or high-level concepts to improve semantic understanding for zero-shot class prediction tasks. Recently, many works [3, 20, 22, 29, 53, 60] also employ LLM to generate fine-grained descriptions for detection tasks. [60] proposes a novel framework that utilizes contextual LLM tokens as conditional object queries to enhance the visual decoder. [53] leverages LLM’s semantic prior as a filter to filter out HOIs that are unlikely to interact. Given the value of external knowledge provided by LLMs, it is crucial to investigate how to efficiently integrate this knowledge for a better understanding of HOI concepts. However, investigating how to integrate this knowledge is

underexplored in the HOI community. In our work, we explore the LLM-generated descriptions tailored for human-centric relationship understanding. Instead of using the general description of the object or action simply, we propose to utilize the body parts descriptions generated by querying LLM for fine-grained semantic enhancement during interaction recognition. The incorporation of body part descriptions facilitates a more nuanced understanding of interactions, enabling greater generalization across diverse human-centric scenarios.

3. Method

In this section, we aim to design an open-vocabulary HOI detector that can detect HOIs described by arbitrary text inputs leveraging large foundation models. An overview of our Open-Vocabulary HOI detector is shown in Figure 2. In the following, we start with the formal problem formulation and introduce a basic pipeline with an encoder-decoder architecture in Section 3.1, harnessing the potential of CLIP [45]. We then introduce the process of multi-level decoding with the conditional matching of our CMD-SE in Section 3.2. This design motivates the model to detect HOIs with different distances with different feature maps, alleviating the potential suboptimal issue. In Section 3.3, we describe the generation and incorporation of fine-grained semantics pertaining to human body parts behind the interaction labels. This approach enhances the model’s ability to discern correlations in human postures across different actions, guiding the model within a more structured semantic space. At last, we present loss functions used to train our CMD-SE in Section 3.4.

3.1. Preliminary

Problem Formulation. We define an interaction as a tuple (b_o, b_h, o, a) , where $b_o, b_h \in \mathbb{R}^4$ denote the detected bounding box of a human and object instance, respectively. $o \in \mathbb{O}$ and $a \in \mathbb{A}$ denote the object and action category, where $\mathbb{A} = \{1, 2, \dots, A\}$ and $\mathbb{O} = \{1, 2, \dots, O\}$ denote the human action and object set, respectively. The objective of open-vocabulary HOI detection is to accurately detect a wide range of interactions. This entails the capability to recognize interactions that have not been encountered during the training phase, encompassing unseen objects, actions, and their various combinations.

A Basic Pipeline. We first build an end-to-end open-vocabulary HOI detector with the help of the generalization capability on classification tasks of CLIP [45] and cast HOI detection as an end-to-end set matching problem similar to DETR [5], eliminating the need for handcrafted components like anchor generation. Given an image \mathbf{I} , the global context representation F_{img} is first extracted by a

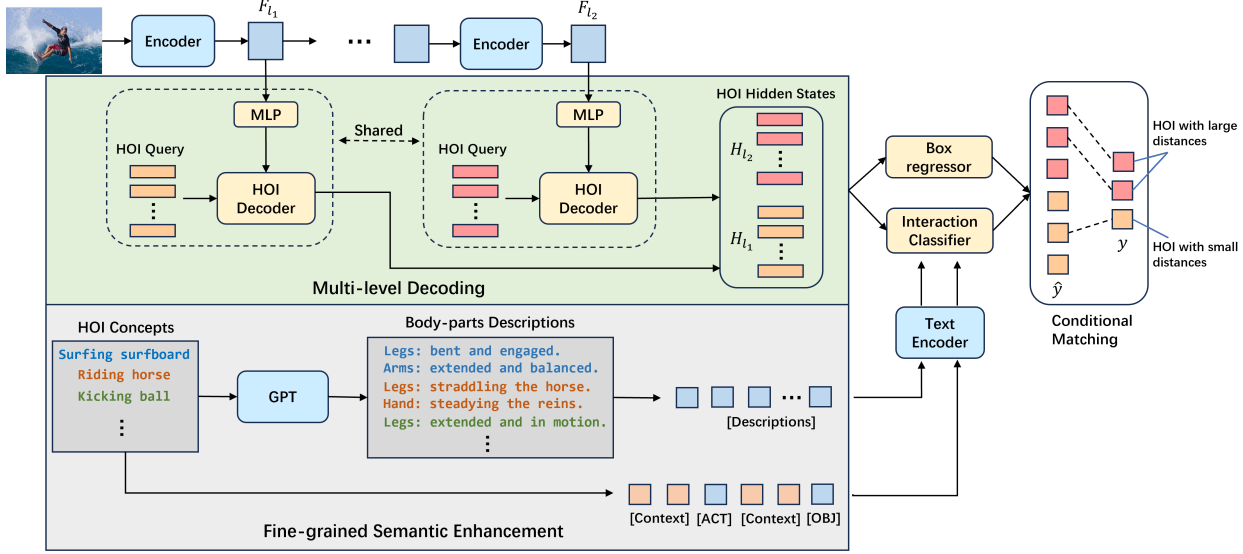


Figure 2. The framework of our CMD-SE. Given an image, the visual encoder is first applied to extract the multi-level visual features. Then we decode the HOIs from multi-level feature maps through a shared HOI decoder parallelly and encourage the HOIs decoded from low-level feature maps to model HOIs with small distances and vice versa via conditional matching. Additionally, we query GPT to describe the states of human body parts for each HOI and utilize the generalizable and fine-grained descriptions as additional prompts to improve interaction recognition.

pre-trained CLIP visual encoder E_V :

$$F_{img} = E_V(\mathbf{I}) \quad (1)$$

where $F_{img} \in \mathbb{R}^{HW \times C}$ denotes a sequence of feature embeddings of \mathbf{I} . We then utilize a transformer-based decoder D to decode the HOIs in \mathbf{I} . Specifically, taking the projected context feature and HOI queries $Q = (q_1, q_2, \dots, q_M)$ as input, the output from the final layer of D serves as the representation of interactions:

$$H = D(Q, F_{img}) \quad (2)$$

where Q is treated as query, and the projected context representation is treated as key and value during the cross-attention mechanism of the HOI decoder D , $H = (h_1, h_2, \dots, h_M)$, where M denotes the number of HOI queries. Then, we feed them to two different head networks: 1) a bounding box regressor P_{bbox} , which predicts a confidence score c and the bounding box of the interacting human and object (b_h, b_o) . 2) a linear projection P_{cls} which maps the feature to the joint vision-and-text space. Similar to [39, 56, 57], we compute its similarity with the features from the text encoder for interaction recognition. Finally we compute the box regression loss \mathcal{L}_b , the intersection-over-union loss \mathcal{L}_{iou} , and the interaction classification loss \mathcal{L}_{cls} . Similar to [34], the matching cost is formulated as:

$$\mathcal{L}_{cost} = \lambda_b \sum_{i \in \{h, o\}} \mathcal{L}_b^i + \lambda_{iou} \sum_{i \in \{h, o\}} \mathcal{L}_{iou}^i + \lambda_{cls} \mathcal{L}_{cls} \quad (3)$$

3.2. Conditional Multi-level Decoding

Since the variance of the distances between interactive human-object pairs becomes larger as the vocabulary space of interaction grows, we propose using multi-level feature maps to model HOIs with different center distances to avoid suboptimal performance [25]. In this subsection, we first formulate the procedure of decoding HOIs from multi-level feature maps. Then we introduce a conditional matching mechanism that encourages different levels of feature maps to model different types of HOIs explicitly, as illustrated in Figure 2.

Multi-level Decoding. The visual encoder E_V takes an image \mathbf{I} with a fixed resolution (e.g., 224×224), divides it into small patches, and first projects them as a sequence. Assume $E_V = (E_V^0, E_V^1, \dots, E_V^N)$, where E_V^i is the i -th block of E_V . Denote the input as X_0 , E_V encodes the image via the self-attention mechanism [54] in each layer:

$$X_{i+1} = E_V^i(X_i) \quad (4)$$

where X_i denotes the encoded feature map after the i -th block of E_V . Then we utilize multi-level feature maps to decode HOIs in parallel:

$$H_{l_i} = D(Q, X_{l_i}), l_i \in \mathbb{L} \quad (5)$$

where \mathbb{L} denotes the set of levels of interest. Then the decoded hidden states of HOIs H_{l_i} ($l_i \in \mathbb{L}$) are fed to the bounding box regressor and the interaction classifier to pre-

dict the bounding boxes of human-object pairs and the interaction categories, respectively.

Conditional Matching. During the bipartite matching [5] between the set of N predictions \hat{y} and the set of ground truth HOIs $y = \{y_i^K\}$ per image, we design a unique loss to explicitly encourage H_i decoded from the low-level feature map to model HOIs with relatively small center distances and vice versa to alleviate the performance gap across HOIs with different distances. Specifically, besides the typical matching loss \mathcal{L}_{cost} introduced in Equation 3, we design a novel constraint loss \mathcal{L}_d between the normalized level index of \hat{y}_i and the H-O distance in y_j :

$$\mathcal{L}_d = \|(Lv(\hat{y}_i), g(y_j))\|_1 \quad (6)$$

where $Lv(\hat{y}_i) \in (0, 1)$ denotes the normalized level index of \hat{y}_i , $g(y_j)$ denotes L_2 distance between the center of the annotated human and object in y_j . This design encourages the model to categorize HOIs with distinct distances into separate groups and address HOIs within each group by utilizing a specific feature map tailored to that group’s characteristics.

3.3. Fine-grained Semantic Enhancement

To enhance the differentiation of various HOI concepts that can be challenging to distinguish based solely on their category names, we propose to leverage the underlying states of human body parts associated with the interaction labels. The descriptions of these states serve to provide a more fine-grained and comprehensive understanding of HOI concepts. In this subsection, we present the process of generating descriptions for the states of human body parts involved in HOIs via prompting with GPT. Additionally, we introduce how the semantics of human body parts are integrated to our model to improve interaction recognition.

Prompt with GPT. To harness the generalizable and recombinant nature of human body parts for interaction recognition, we employ GPT [2] to produce state descriptions of body parts. However, directly querying GPT for body part descriptions may yield verbose descriptions, neglecting the varying relevance of different body parts in interactions. Consequently, we implement a two-step mechanism to query GPT for generating body part descriptions for each HOI. As illustrated in Figure 3, our approach involves utilizing a standardized pipeline and set of templates as language commands to interact with GPT for each HOI category. Specifically, given an HOI name consisting of an action name act_i , and an object name obj_i , we initiate the query to GPT by asking: “Which body parts are involved in the interaction when a person act_i a/an obj_i ?” This command aids in identifying a subset of body parts BP_i that are pertinent to the current i -th interaction:

$$BP_i = GPT(BP, Query_i) \quad (7)$$

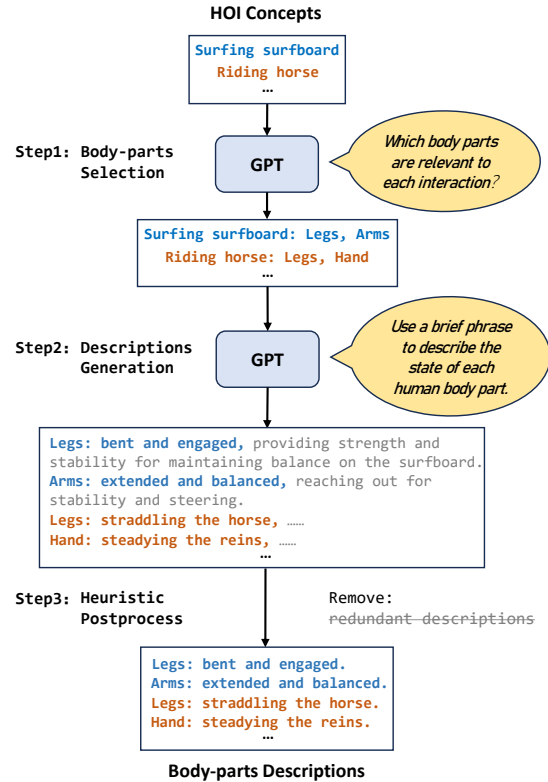


Figure 3. Illustration of generating prompts with GPT. The main purpose of the entire process is to find the most simple and general body parts descriptions for each HOI.

where BP represents a predefined list of human body parts, including “mouth, eye, arm, hand, leg, foot”. For example, if the interaction is “surfing surfboard”, the returned BP_i from GPT would include arm and leg. The main purpose of this step is to limit the description to a few important body parts. If it is generated directly without adding this restriction, the description will not be able to distinguish between key body parts and non-key body parts. This will make it difficult for us to accurately identify the characteristics of the action state and cause confusion in HOI concepts.

Then we further query GPT with the prompt $Query'_i$: “Use a brief phrase to describe the state of each human body part in BP_i when a person act_i a/an obj_i .” to acquire the state descriptions S_i of the corresponding human body parts:

$$S_i = GPT(BP_i, Query'_i) \quad (8)$$

Finally, we use a heuristic postprocess to remove some redundant descriptions, making the status description more simple and general.

Enhance Interaction Recognition. The primary objective of interaction recognition is to convert the textual description of interactions into a feature space and compare them with the output generated by the HOI decoder in Sec-

tion 3.2. In order to address the limited distinctions among HOI names, we also propose leveraging the semantic information provided by descriptions of body parts associated with each HOI. This means that for each HOI, we encode not only its name but also the states of the relevant human body parts. Following [56, 69], given a HOI category defined as an action-object pair, denoted as (act_i, obj_i) , we form a sentence using global learnable context tokens Ctx . As illustrated in Figure 2, we use a few prefix tokens Ctx_{pre} at the beginning of the sentence and a few learnable conjunct tokens Ctx_{con} to automatically determine how to connect the category names of the action and object. Then, we generate the text embedding for each HOI label through the pre-trained CLIP text-encoder E_T :

$$T_{hoi} = E_T(Ctx_{pre}, act_i, Ctx_{con}, obj_i) \quad (9)$$

The resulting embeddings of all HOI names are denoted as $T_{hoi} \in \mathbb{R}^{N \times C_t}$, where N denotes the number of HOI triplet categories, and C_t represents the dimension of text embedding from CLIP text encoder.

Furthermore, as the descriptions of body parts are already semantically rich, they do not require learnable vectors as context. To encode the state descriptions S_i of the human body parts for each HOI, as obtained in Equation 8, we directly utilize the CLIP text encoder E_T :

$$T_b = E_T(S_i) \quad (10)$$

where $T_b \in \mathbb{R}^{N \times C_t}$ represents the embeddings of body parts' descriptions for all HOIs. Note that we only need to forward the body parts descriptions once since there are no learnable parameters in this branch, which significantly enhances its efficiency. The final interaction prediction is obtained by combining the output logits of T_{hoi} and T_b :

$$p_j = \alpha_{hoi} h_j T_{hoi}^T + \alpha_b h_j T_b^T \quad (11)$$

where h_j represents the hidden state of a HOI, p_j represents the corresponding interaction prediction, α_{hoi} and α_b are learnable scalars used to weigh the contributions of the HOI embeddings (T_{hoi}) and body part embeddings (T_b) respectively.

3.4. Training and Inference

In this subsection, we elaborate the processes of training and inference of our model.

Training. During the training stage, we follow the query-based methods [5, 34, 50, 56] to assign a bipartite matching prediction with each ground-truth using the Hungarian algorithm [27]. The matching cost \mathcal{L}_{cost} for the matching process and the targeting cost for the training back-propagation share the same strategy, which is formulated in Section 3.1. Then considering the additional constraint loss \mathcal{L}_d introduced in Section 3.2, the final matching loss can be defined

as:

$$\mathcal{L} = \mathcal{L}_{cost} + \lambda_d \mathcal{L}_d \quad (12)$$

where λ_d is the hyper-parameter weight for the additional constraint.

Inference. For each HOI prediction, including the bounding-box pair $(\hat{b}_h^i, \hat{b}_o^i)$, the bounding box score \hat{c}_i from the box regressor, and the interaction score \hat{s}_i from the interaction classifier, the final score \hat{s}_i' is computed as:

$$\hat{s}_i' = \hat{s}_i \cdot \hat{c}_i^\gamma \quad (13)$$

where $\gamma > 1$ is a constant used during inference to suppress overconfident objects [62, 63].

4. Experiment

4.1. Experimental Setting

Datasets. Our experiments are mainly conducted on two datasets, SWIG-HOI [55] and HICO-DET [6]. The SWIG-HOI dataset encompasses 400 human actions and 1000 object categories. The test set of SWIG-HOI comprises approximately 14,000 images and 5,500 interactions, with around 1,800 interactions not present in the training set. The HICO-DET dataset provides 600 combinations involving 117 human actions and 80 objects. We follow [16, 56] to simulate a zero-shot detection setting by holding out 120 rare interactions from all 600 interactions.

Evaluation Metric. We follow the settings of previous works [6, 34, 35, 56] to use the mean Average Precision (mAP) for evaluation. We define an HOI triplet prediction as a true-positive example if the following criteria are met: 1) The IoU of the human bounding box and object bounding box are larger than 0.5 *w.r.t.* the GT bounding boxes; 2) the predicted interaction category is accurate.

Implementation Details. Our model is built upon the pre-trained CLIP and all its parameters are frozen during training. We employ the ViT-B/16 version as our visual encoder following [56]. We set the cost weights λ_b , λ_{iou} , λ_{cls} and λ_d to 5, 2, 5, and 5 during training and train our model for 80 epochs with a batch size of 128 on 2 A100 GPUs.

4.2. Comparison with Other Methods

We compare the performance of our model with existing methods on SWIG-HOI and HICO-DET datasets. As shown in Table 1, our model significantly outperforms the previous state-of-the-art on the SWIG-HOI dataset on all splits, achieving a relative gain of 15.08% on all interactions. This shows the strong capability of our CMD-SE to detect and recognize the interactions of human-object pairs in the open-vocabulary scenario.

We also compare our method with state-of-the-art methods on the simulated zero-shot setting of the HICO-DET dataset in Table 2. Note that earlier methods [16, 18] lack

Method	Non-rare	Rare	Unseen	Full
QPIC [50]	16.95	10.84	6.21	11.12
THID [56]	17.67	12.82	10.04	13.26
CMD-SE (Ours)	21.46	14.64	10.70	15.26

Table 1. Comparison of our proposed CMD-SE with state-of-the-art methods on the SWIG-HOI dataset.

Method	Pretrained Detector	Unseen	Seen	Full
<i>Zero-shot Methods</i>				
VCL [16]	✓	10.06	24.28	21.43
ATL [17]	✓	9.18	24.67	21.57
FCL [18]	✓	13.16	24.23	22.01
GEN-VLKT [34]	✓	21.36	32.91	30.56
HOICLIP [39]	✓	23.48	34.47	32.26
<i>Open-vocabulary Methods</i>				
THID [56]	✗	15.53	24.32	22.38
CMD-SE (Ours)	✗	16.70	23.95	22.35

Table 2. Comparison of our proposed CMD-SE with state-of-the-art methods on HICO-DET under the simulated zero-shot setting.

Method	Non-rare	Rare	Unseen	Full
<i>Base</i>	15.69	11.53	7.32	11.45
+ <i>MD</i>	18.66	12.93	7.29	12.87
+ <i>CM</i>	21.33	14.26	9.32	14.69
+ <i>SE</i>	21.46	14.64	10.70	15.26

Table 3. Ablations of different modules of our CMD-SE on the SWIG-HOI dataset. *MD*: multi-level decoding. *CM*: conditional matching. *SE*: semantic enhancement.

semantic understanding of interaction categories and pre-determine the unseen interactions during the training stage. Although recent zero-shot methods [34, 39] employ CLIP text embeddings for interaction classification, they typically rely on a DETR architecture and pretrained DETR weights (pretrained on object detection datasets with finite categories), which is unscalable in terms of the vocabulary size, thereby limiting their applicability in open-world scenarios. Note that it’s not a fair comparison between our method and theirs on HICO-DET, as the HICO-DET dataset and the COCO dataset used for DETR pretraining share the same object label space. In contrast, open-vocabulary methods [56] break free from this constraint and do not pretrain on any detection datasets. In comparison to the previous open-vocabulary HOI detector, our CMD-SE achieves new state-of-the-art performance on unseen HOIs.

4.3. Ablation Study

Network Architecture. In this section, we empirically investigate how the performance of the proposed method is affected by different model settings on the SWIG-HOI dataset. We denote the basic pipeline model introduced in Section 3.1 as *Base* in Table 3. The first modification is applying multi-level features for HOI decoding. We observe an improvement of 1.42% mAP. Subsequently, we apply the soft constraint to perform conditional matching, denoted as +*CM*, explicitly guiding different levels of feature maps to represent HOIs with varying distances. The results show that +*CM* brings a significant improvement across all splits, contributing to a 2.4% mAP gain on all categories. To further enhance the model’s semantic understanding of interaction concepts, we incorporate the fine-grained semantics of human body parts, denoted as +*SE*. This modification yields the best performance, with a notable relative gain of 22.8% in the mAP for unseen classes. This outcome underscores the substantial potential of our framework in enhancing interaction understanding within an open-vocabulary scenario.

Distance Types and Matching Strategies. In this section, we explore the impact of different H-O distance types and matching strategies on model performance. Initially, we examine the effectiveness of absolute versus relative distances for matching. While intuition may suggest relative distance would be more effective, as it encapsulates inherent HOI properties, our results, depicted in the first two lines of Table 4, empirically favor absolute distance. We attribute this to its alignment with ViT’s attention mechanism, where each patch’s attention expands outward by a distance corresponding to the absolute distance in the image. As shown in Table 4, matching low-level features with short-distance HOIs achieves superior results compared to matching low-level features for long-distance HOIs, with an improvement of 0.45 and 1.33 mAP on seen and unseen classes, respectively. This finding suggests that low-level features are more effective in capturing interactions that occur in relatively close proximity.

Additional Constraint for Conditional Matching. We discuss the weight of the additional constraint during bipartite matching. As shown in Table 5, when setting the weight to 0, the model is optimized to detect HOIs with different distances with the same set of feature maps, leading to suboptimal performance due to the variation in distances between interactive humans and objects. Setting the weight to 5 encourages our model to differentiate and model HOIs with different distances using distinct feature maps, leading to the best performance and a notable 1.46 mAP gain. However, increasing the weight further to 10 causes a marginal decline in the model’s performance. This can be attributed to the higher weight potentially causing the model to neglect other pertinent factors during bipartite matching.



Figure 4. Qualitative results of our method on SWIG-HOI test set.

Distance Types	Matching Strategies	Seen	Unseen
Relative	Low-Small	16.57	10.30
Absolute	Low-Small	16.79	10.70
Absolute	Low-Large	16.61	10.47

Table 4. Ablations on the H-O distance types and matching strategies between the levels of feature maps and the H-O distances. *Relative*: using relative distance as H-O distance. *Absolute*: using absolute distance as H-O distance. *Low-Small*: matching low layer with small distance HOI. *Low-Large*: matching low layer with large distance HOI.

λ_d	Non-rare	Rare	Unseen	Full
0	19.80	13.94	7.89	13.80
5	21.46	14.64	10.70	15.26
10	20.18	14.06	9.64	14.40

Table 5. Ablations of the weight for the additional soft constraint during conditional matching.

Prompts for Fine-grained Semantic Enhancement. In this part, we explore different types of prompts for enhancing fine-grained semantics. A naive and straightforward approach is to describe the action and object associated with each HOI, termed as ‘‘HOI Descriptions’’. As shown in the first line of Table 6, utilizing HOI descriptions leads to a mAP of 16.92 and 10.52 on seen and unseen classes, respectively. Considering the generalizable and recombinant nature of human body parts, we further explore employing body parts for semantic enhancement. Simply utilizing the names of body parts relevant to each HOI, generated by GPT following the ‘‘Body-part Selection’’ procedure in Figure 3, results in a marginally inferior performance. This is attributed to the finite set of body part combinations and the presence of the same set of body part names across different HOIs, potentially causing confusion in the model’s ability to distinguish them. By further incorporating descriptions of the states of body parts—characteristics that are not only generalizable and recombinant but also specific to each HOI, our model achieves the best performance.

4.4. Qualitative Results

We visualize the prediction result and attention maps to illustrate the characteristics of our method. As shown in Fig-

Prompts	Seen	Unseen
HOI Descriptions	15.73	10.10
Body Parts Names	16.08	9.81
Body Parts Descriptions	16.79	10.70

Table 6. Ablations on different types of prompts for fine-grained semantic enhancement.

ure 4a and 4b, we find that our model can adaptively use the output of different layers to predict HOIs with different human-object pairs distances. For a farther interaction of communicating person and a closer interaction of kissing person, their optimal predictions are encoded by the output of the 8th layer and the 5th layer of the encoder respectively. This shows that our model has generalization ability in different human-object pairs distances. Additionally, as shown in Figure 4c and 4d, we find that our model exhibits some sensitivity to human body parts. For these two examples, actions of hurling and kicking, our CMD-SE focuses on the hands and arms for hurling, and the legs for kicking. This shows that our model can use the characteristics of human body parts to infer the HOI interaction through our design.

5. Conclusion

We utilize large foundation models to build an end-to-end open vocabulary HOI detector that addresses two major challenges in the field. First, we tackle the issue of diverse spatial distances between interactive human-object pairs by utilizing different levels of feature maps tailored to HOIs at different distances. Second, we leverage fine-grained semantics of human body parts, in addition to category names, to enhance interaction recognition. By querying a LLM for descriptions of human body parts given a HOI, the detector gains a deeper understanding of the correlation of human postures between different actions. Experimental results demonstrate the superior performance of our model in open-vocabulary HOI detection.

Acknowledgements. This work was supported by the grants from the National Natural Science Foundation of China 62372014.

References

- [1] Ankan Bansal, Sai Saketh Rambhatla, Abhinav Shrivastava, and Rama Chellappa. Detecting human-object interactions via functional generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10460–10469, 2020. [2](#)
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [2](#), [3](#), [5](#)
- [3] Yichao Cao, Qingfei Tang, Xiu Su, Chen Song, Shan You, Xiaobo Lu, and Chang Xu. Detecting any human-object interaction relationship: Universal hoi detector with spatial prompt learning on foundation models, 2023. [3](#)
- [4] Yichao Cao, Qingfei Tang, Feng Yang, Xiu Su, Shan You, Xiaobo Lu, and Chang Xu. Re-mine, learn and reason: Exploring the cross-modal semantic correlations for language-guided hoi detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23492–23503, 2023. [1](#), [2](#)
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. [3](#), [5](#), [6](#)
- [6] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *2018 IEEE winter conference on applications of computer vision (wacv)*, pages 381–389. IEEE, 2018. [6](#)
- [7] Mingfei Chen, Yue Liao, Si Liu, Zhiyuan Chen, Fei Wang, and Chen Qian. Reformulating hoi detection as adaptive set prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9004–9013, 2021. [2](#)
- [8] Leizhen Dong, Zhimin Li, Kunlun Xu, Zhijun Zhang, Luxin Yan, Sheng Zhong, and Xu Zou. Category-aware transformer network for better human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19538–19547, 2022. [1](#)
- [9] Mohamed Elhoseiny, Yizhe Zhu, Han Zhang, and Ahmed Elgammal. Link the head to the” beak”: Zero shot learning from noisy text description at part precision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5640–5649, 2017. [3](#)
- [10] Christiane Fellbaum. *WordNet: An electronic lexical database*. MIT press, 1998. [3](#)
- [11] Chen Gao, Yuliang Zou, and Jia-Bin Huang. ican: Instance-centric attention network for human-object interaction detection. *arXiv preprint arXiv:1808.10437*, 2018. [2](#)
- [12] 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. [2](#)
- [13] 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. [2](#)
- [14] Tanmay Gupta, Alexander Schwing, and Derek Hoiem. No-frills human-object interaction detection: Factorization, layout encodings, and training techniques. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9677–9685, 2019. [2](#)
- [15] Xiangteng He and Yuxin Peng. Fine-grained image classification via combining vision and language. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5994–6002, 2017. [3](#)
- [16] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. In *European Conference on Computer Vision*, pages 584–600. Springer, 2020. [2](#), [6](#), [7](#)
- [17] Zhi Hou, Baosheng Yu, Yu Qiao, Xiaojiang Peng, and Dacheng Tao. Affordance transfer learning for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 495–504, 2021. [7](#)
- [18] Zhi Hou, Baosheng Yu, Yu Qiao, Xiaojiang Peng, and Dacheng Tao. Detecting human-object interaction via fabricated compositional learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14646–14655, 2021. [2](#), [6](#), [7](#)
- [19] ASM Iftekhar, Hao Chen, Kaustav Kundu, Xinyu Li, Joseph Tighe, and Davide Modolo. What to look at and where: Semantic and spatial refined transformer for detecting human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5353–5363, 2022. [3](#)
- [20] Sheng Jin, Xueying Jiang, Jiaying Huang, Lewei Lu, and Shijian Lu. Llm meets vlm: Boost open vocabulary object detection with fine-grained descriptors, 2024. [3](#)
- [21] Keizo Kato, Yin Li, and Abhinav Gupta. Compositional learning for human object interaction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–251, 2018. [2](#)
- [22] Prannay Kaul, Weidi Xie, and Andrew Zisserman. Multi-modal classifiers for open-vocabulary object detection, 2023. [3](#)
- [23] Bumsoo Kim, Taeho Choi, Jaewoo Kang, and Hyunwoo J Kim. Uniondet: Union-level detector towards real-time human-object interaction detection. In *European Conference on Computer Vision*, pages 498–514. Springer, 2020. [2](#)
- [24] Bumsoo Kim, Junhyun Lee, Jaewoo Kang, Eun-Sol Kim, and Hyunwoo J Kim. Hotr: End-to-end human-object interaction detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 74–83, 2021. [1](#), [2](#)
- [25] Bumsoo Kim, Jonghwan Mun, Kyoung-Woon On, Minchul Shin, Junhyun Lee, and Eun-Sol Kim. Mstr: Multi-scale transformer for end-to-end human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19578–19587, 2022. [2](#), [3](#), [4](#)
- [26] Sanghyun Kim, Deunsol Jung, and Minsu Cho. Relational context learning for human-object interaction detection. In

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2925–2934, 2023. 2
- [27] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 2, 6
- [28] Ting Lei, Fabian Caba, Qingchao Chen, Hailin Ji, Yuxin Peng, and Yang Liu. Efficient adaptive human-object interaction detection with concept-guided memory. 2023. 2
- [29] Lin Li, Jun Xiao, Guikun Chen, Jian Shao, Yueting Zhuang, and Long Chen. Zero-shot visual relation detection via composite visual cues from large language models, 2023. 3
- [30] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactiveness knowledge for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019. 2
- [31] 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
- [32] Zhimin Li, Cheng Zou, Yu Zhao, Boxun Li, and Sheng Zhong. Improving human-object interaction detection via phrase learning and label composition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1509–1517, 2022. 1
- [33] 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. 2
- [34] Yue Liao, Aixi Zhang, Miao Lu, Yongliang Wang, Xiaobo Li, and Si Liu. Gen-vlkt: Simplify association and enhance interaction understanding for hoi detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20123–20132, 2022. 2, 3, 4, 6, 7
- [35] Xinpeng Liu, Yong-Lu Li, Xiaoqian Wu, Yu-Wing Tai, Cewu Lu, and Chi-Keung Tang. Interactiveness field in human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20113–20122, 2022. 1, 2, 6
- [36] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. *arXiv preprint arXiv:2210.07183*, 2022. 3
- [37] Muhammad Ferjad Naeem, Yongqin Xian, Luc V Gool, and Federico Tombari. I2dformer: Learning image to document attention for zero-shot image classification. *Advances in Neural Information Processing Systems*, 35:12283–12294, 2022. 3
- [38] Muhammad Ferjad Naeem, Muhammad Gul Zain Ali Khan, Yongqin Xian, Muhammad Zeshan Afzal, Didier Stricker, Luc Van Gool, and Federico Tombari. I2mvformer: Large language model generated multi-view document supervision for zero-shot image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15169–15179, 2023. 3
- [39] Shan Ning, Longtian Qiu, Yongfei Liu, and Xuming He. Hoiclip: Efficient knowledge transfer for hoi detection with vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23507–23517, 2023. 2, 3, 4, 7
- [40] Zachary Novack, Julian McAuley, Zachary Chase Lipton, and Saurabh Garg. Chils: Zero-shot image classification with hierarchical label sets. In *International Conference on Machine Learning*, pages 26342–26362. PMLR, 2023. 3
- [41] JeeSeung Park, Jin-Woo Park, and Jong-Seok Lee. Viplo: Vision transformer based pose-conditioned self-loop graph for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17152–17162, 2023. 2
- [42] Sarah Pratt, Ian Covert, Rosanne Liu, and Ali Farhadi. What does a platypus look like? generating customized prompts for zero-shot image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15691–15701, 2023. 3
- [43] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2
- [44] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 2
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 3
- [46] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 49–58, 2016. 3
- [47] Karsten Roth, Oriol Vinyals, and Zeynep Akata. Integrating language guidance into vision-based deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16177–16189, 2022. 3
- [48] Karsten Roth, Jae Myung Kim, A. Sophia Koepke, Oriol Vinyals, Cordelia Schmid, and Zeynep Akata. Waffling around for performance: Visual classification with random words and broad concepts, 2023. 3
- [49] Sheng Shen, Chunyuan Li, Xiaowei Hu, Yujia Xie, Jianwei Yang, Pengchuan Zhang, Zhe Gan, Lijuan Wang, Lu Yuan, Ce Liu, et al. K-lite: Learning transferable visual models with external knowledge. *Advances in Neural Information Processing Systems*, 35:15558–15573, 2022. 3
- [50] Masato Tamura, Hiroki Ohashi, and Tomoaki Yoshinaga. Qpic: Query-based pairwise human-object interaction detection with image-wide contextual information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10410–10419, 2021. 2, 6, 7

- [51] Danyang Tu, Wei Sun, Guangtao Zhai, and Wei Shen. Agglomerative transformer for human-object interaction detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21614–21624, 2023. [2](#)
- [52] Oytun Ulutan, ASM Iftekhar, and Bangalore S Manjunath. Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13617–13626, 2020. [2](#)
- [53] Mesut Erhan Unal and Adriana Kovashka. Weakly-supervised hoi detection from interaction labels only and language/vision-language priors, 2023. [3](#)
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [4](#)
- [55] Suchen Wang, Kim-Hui Yap, Henghui Ding, Jiyan Wu, Junsong Yuan, and Yap-Peng Tan. Discovering human interactions with large-vocabulary objects via query and multi-scale detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13475–13484, 2021. [3](#), [6](#)
- [56] Suchen Wang, Yueqi Duan, Henghui Ding, Yap-Peng Tan, Kim-Hui Yap, and Junsong Yuan. Learning transferable human-object interaction detectors with natural language supervision. In *CVPR*, 2022. [2](#), [3](#), [4](#), [6](#), [7](#)
- [57] Mingrui Wu, Jiaxin Gu, Yunhang Shen, Mingbao Lin, Chao Chen, Xiaoshuai Sun, and Rongrong Ji. End-to-end zero-shot hoi detection via vision and language knowledge distillation. *arXiv preprint arXiv:2204.03541*, 2022. [2](#), [3](#), [4](#)
- [58] Chi Xie, Fangao Zeng, Yue Hu, Shuang Liang, and Yichen Wei. Category query learning for human-object interaction classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15275–15284, 2023. [1](#)
- [59] Dongming Yang and Yuexian Zou. A graph-based interactive reasoning for human-object interaction detection. *arXiv preprint arXiv:2007.06925*, 2020. [2](#)
- [60] Yuhang Zang, Wei Li, Jun Han, Kaiyang Zhou, and Chen Change Loy. Contextual object detection with multimodal large language models, 2023. [3](#)
- [61] Aixi Zhang, Yue Liao, Si Liu, Miao Lu, Yongliang Wang, Chen Gao, and Xiaobo Li. Mining the benefits of two-stage and one-stage hoi detection. *Advances in Neural Information Processing Systems*, 34:17209–17220, 2021. [1](#), [2](#)
- [62] Frederic Z. Zhang, Dylan Campbell, and Stephen Gould. Spatially conditioned graphs for detecting human-object interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13319–13327, 2021. [6](#)
- [63] Frederic Z Zhang, Dylan Campbell, and Stephen Gould. Efficient two-stage detection of human-object interactions with a novel unary-pairwise transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20104–20112, 2022. [1](#), [2](#), [6](#)
- [64] Frederic Z. Zhang, Yuhui Yuan, Dylan Campbell, Zhuoyao Zhong, and Stephen Gould. Exploring predicate visual context in detecting human-object interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10411–10421, 2023. [2](#)
- [65] Yong Zhang, Yingwei Pan, Ting Yao, Rui Huang, Tao Mei, and Chang-Wen Chen. Exploring structure-aware transformer over interaction proposals for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19548–19557, 2022. [1](#)
- [66] Sipeng Zheng, Boshen Xu, and Qin Jin. Open-category human-object interaction pre-training via language modeling framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19392–19402, 2023. [3](#)
- [67] Xubin Zhong, Changxing Ding, Xian Qu, and Dacheng Tao. 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*, pages 69–85. Springer, 2020. [3](#)
- [68] Xubin Zhong, Xian Qu, Changxing Ding, and Dacheng Tao. 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*, pages 13234–13243, 2021. [2](#)
- [69] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [6](#)
- [70] Penghao Zhou and Mingmin Chi. Relation parsing neural network for human-object interaction detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 843–851, 2019. [3](#)