

Open-World Human-Object Interaction Detection via Multi-modal Prompts

Jie Yang^{1,2*}, Bingliang Li^{1*}, Ailing Zeng^{2†}, Lei Zhang², Ruimao Zhang^{1†}

¹The Chinese University of Hong Kong, Shenzhen ²International Digital Economy Academy

Abstract

In this paper, we develop **MP-HOI**, a powerful **Multi-modal Prompt-based HOI** detector designed to leverage both textual descriptions for open-set generalization and visual exemplars for handling high ambiguity in descriptions, realizing HOI detection in the open world. Specifically, it integrates visual prompts into existing language-guided-only HOI detectors to handle situations where textual descriptions face difficulties in generalization and to address complex scenarios with high interaction ambiguity. To facilitate MP-HOI training, we build a large-scale HOI dataset named **Magic-HOI**, which gathers six existing datasets into a unified label space, forming over 186K images with 2.4K objects, 1.2K actions, and 20K HOI interactions. Furthermore, to tackle the long-tail issue within the **Magic-HOI** dataset, we introduce an automated pipeline for generating realistically annotated HOI images and present **SynHOI**, a high-quality synthetic HOI dataset containing 100K images. Leveraging these two datasets, MP-HOI optimizes the HOI task as a similarity learning process between multi-modal prompts and objects/interactions via a unified contrastive loss, to learn generalizable and transferable objects/interactions representations from large-scale data. MP-HOI could serve as a generalist HOI detector, surpassing the HOI vocabulary of existing expert models by more than 30 times. Concurrently, our results demonstrate that MP-HOI exhibits remarkable zero-shot capability in real-world scenarios and consistently achieves a new state-of-the-art performance across various benchmarks. Our project homepage is available at <https://MP-HOI.github.io/>.

1. Introduction

Human-Object Interaction Detection (HOI), a core element in human-centric vision perception tasks such as human activity recognition [15], motion tracking [58], and anomaly behavior detection [31], has attracted considerable attention over the past decades. HOI primarily involves localizing


correlated human-object 2D positions within an image and identifying their interactions. Although numerous models have been proposed [17, 30, 50, 55, 61, 66], deploying these models in open-world scenarios remains a significant challenge due to their primary training on closed-set data and limited generalization capabilities.

With the advancements in visual-linguistic models, such as CLIP [42], recent research [30, 36, 54] attempts to introduce the natural language into HOI task and transfer knowledge from CLIP to recognize unseen HOI concepts. However, these methods still suffer from the following limitations: (1) **Limited data and category definitions**: They primarily rely on training data from a single dataset, such as HICO-DET [7], which is built upon a relatively small set of predefined objects (e.g., 80) and actions (e.g., 117). Consequently, this narrow range of object and action categories restricts their ability to generalize effectively. (2) **An inherent bottleneck with textual descriptions**: Even when the model is expanded to be trained on more objects and action categories, it still encounters a bottleneck when dealing with entirely new categories that lack any relevant text descriptions in the training data. (3) **Difficulty addressing high-ambiguity scenarios**: Current methods typically associate a human with a specific object based solely on a single interaction description. However, real-world situations often involve the composition of multiple interactions, as depicted in Fig 1-(a), introducing significant ambiguity in interaction learning. Moreover, these models may fail to detect a person engaged in a specific combination of multiple HOIs, which is essential for certain practical applications.

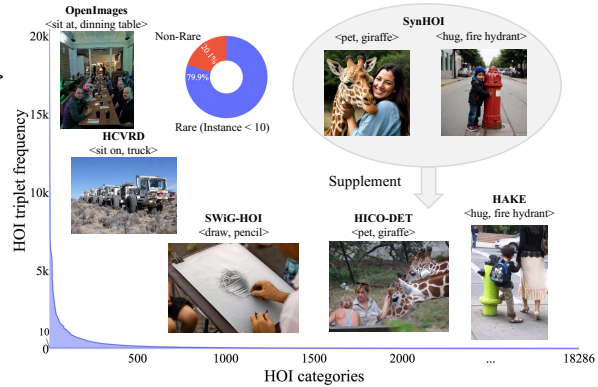
To tackle the above issues, we develop **MP-HOI**, a novel and powerful Multi-modal Prompt-based HOI generalist model for effective HOI detection in the open world. Within **MP-HOI**, we incorporate the visual prompts into existing language-guided-only HOI detectors to enhance the effectiveness and generalization capability. Visual prompts can act as exemplars, allowing the model to detect the same objects/interactions (single HOI or the combination of multiple HOIs for the same person), and they also aid in reducing semantic ambiguity present in textual prompts. Simultaneously, textual prompts also provide semantic context to enhance the understanding of visual prompts. The proposed

*Equal contribution.

†Corresponding author.



(a) Composition of multiple HOIs in the open world



(b) The Distribution of HOI Triplets on Magic-HOI

Figure 1. We show (a) the coexisting composited interactions within the same person in an in-the-wild (e.g., A man is squatting on the ground, holding a painting brush, and painting the wall); (b) the long-tail distribution issue in our *Magic-HOI* dataset, along with another proposed *SynHOI* dataset to address it.

MP-HOI optimize the HOI task as a similarity learning process between multi-modal prompts and objects/interactions by using a unified cross-modality contrastive loss. It allows the model to learn the textual/visual prompts-to-objects/interactions alignments from large-scale data.

To facilitate the effective training of our model, we build a large-scale HOI dataset named *Magic-HOI*, which gathers six existing datasets into a unified label space, forming the one containing over 186K images with 2.4K objects, 1.2K actions, and 20K HOI interactions. Furthermore, as shown in Fig. 1-(b), we analyze the distribution of all HOI interactions in *Magic-HOI*, finding a significant long-tail distribution. Upon further investigation, we discover that 79.9% of HOI interactions have fewer than 10 instances, and over 5900 HOI categories have only a single image available. Such imbalanced label distribution challenges prompt-based learning, particularly visual prompts. Unfortunately, our further investigation reveals that rare HOI categories are genuinely rare in real-life scenarios, making it difficult to collect data from the internet.

Inspired by the impressive conditioned image generation capabilities demonstrated by recent text-to-image diffusion models [1, 43–45, 67], we present a diffusion-based automatic pipeline, including the HOIPrompt design, automatic labeling and filtering, and quality verification, designed to systematically scale up the generation of diverse and high-precision HOI-annotated data. Based on it, we present *SynHOI*, a high-diversity synthetic HOI dataset with over 100K fully annotated HOI images, which can fill in the shortcomings of the long-tail *MP-HOI* dataset.

By leveraging the newly collected *MP-HOI* and *SynHOI* datasets, we train the *MP-HOI* to be a generalist HOI detector, surpassing the HOI vocabulary of existing expert models (e.g., GEN-VLKT [30]) by more than 30 times. Through extensive experiments, we demonstrate that *MP-HOI* exhibits remarkable zero-shot capability in real-world

scenarios. Furthermore, *MP-HOI* consistently achieves the new state-of-the-art performance across various benchmarks. The main contributions of this paper are threefold.

(1) We develop the first multi-modal prompts-based generalist HOI detector, called *MP-HOI*, which leverages both textual descriptions for open-set generalization and visual exemplars for handling high ambiguity in descriptions, paving a new path to HOI detection in the open world.

(2) We build a large-scale *Magic-HOI* dataset for effective large-scale training, which gathers six existing datasets into a unified semantic space. To address the long-tail issue in *Magic-HOI*, we construct an automated pipeline for generating realistically annotated HOI images and present *SynHOI*, a high-quality synthetic HOI dataset.

(3) Extensive experimental results demonstrate that the proposed *MP-HOI* exhibits remarkable zero-shot capability in real-world scenarios and consistently achieves a new state-of-the-art performance across various benchmarks.

2. Related Work

HOI Detection. HOI detection task primarily encompasses three sub-problems, *i.e.*, object detection, human-object pairing, and interaction recognition. Previous HOI detectors can generally be divided into one-stage and two-stage paradigms. The two-stage strategy employs an off-the-shelf detector to determine the locations and classes of objects, followed by specially-designed modules for human-object association and interaction recognition. Most methods are dedicated to exploring additional feature streams to improve interaction classification, such as the appearance stream [11, 19, 20, 25], spatial stream [2, 26, 56], pose and body-parts stream [14, 26, 50], semantic stream [2, 12, 33], and graph network [37, 40, 51, 56, 62]. Instead, the one-stage strategy detects HOI triplets in a single forward pass by assigning human and object proposals to predefined anchors and then estimating potential interactions [10, 20, 29, 53].

Datasets	Objects	Interactions	HOIs	Unified Images	Unified Instances
HICO-DET [7]	80	117	600	37,633	117,871
SWiG-HOI [52]	1,000	407	13,520	50,861	73,499
OpenImages [23]	80	117	600	20,140	78,792
PIC [39]	80	117	600	2,736	9,503
HCVRD [68]	1,821	884	6,311	41,586	256,550
HAKI [24]	80	117	600	33,059	91,120
<i>Magic-HOI</i>	2,427	1,227	20,036	186,015	627,335

Table 1. Statistics of *Magic-HOI* that unifies six existing datasets. Displayed numerically within the table are the counts of attributes we have integrated from each source dataset.

Recently, the DETR-based HOI detectors [8, 21, 35, 47, 48] have gained prominence in this paradigm, which formulate the HOI detection task as a set prediction problem, avoiding complex post-processing. In particular, many methods [4, 9, 22, 30, 36, 59–61, 65, 66] demonstrate promising performance improvements by disentangling human-object detection and interaction classification as two decoders in a cascade manner. Our work builds on the transformer-based HOI detection strategy and focuses on developing the first multi-modal prompts-based generalist HOI detector.

Zero-shot HOI Detection. Zero-shot HOI detection has emerged as a field aiming to identify unseen HOI triplet categories not present in the training data. Previous research [2, 14, 17–19, 33, 38, 38] has addressed this task in a compositional manner, by disentangling the reasoning process on actions and objects during training. This approach enables the recognition of unseen HOI triplets during inference. With the advancements in Vision-Language Models, such as CLIP [42], recent research [30, 36, 54] has shifted focus toward transferring knowledge from CLIP to recognize unseen HOI concepts. This shift has resulted in a notable performance improvement in zero-shot settings. Our work aims to develop a HOI detector further to generalize in open-world scenes effectively.

3. Datasets

This section introduces the proposed two datasets: 1) *Magic-HOI*: A unified large-scale HOI dataset that gathers six existing datasets to form a unified one with consistent label space; 2) *SynHOI*: A high-quality synthetic HOI dataset to tackle the long-tail issue within *Magic-HOI*.

3.1. Magic-HOI

Existing HOI detection datasets have incorporated a wide range of object and interaction categories. However, current models are typically trained on separate datasets, resulting in limited generalization capabilities. In light of this, as shown in Tab. 1, we propose *Magic-HOI* by unifying six HOI detection datasets, including HICO-DET [7], SWiG-HOI [52], OpenImages [23], PIC [39], HCVRD [68], HAKI [24]. To ensure the integrity of *Magic-HOI*, we meticulously curate the object and verb labels across these datasets, diligently addressing issues of label overlap and potential data leakage for zero-shot testing. *Magic-HOI*

Human: "a/an (asian, black, hispanic, ...) (boy, elderly, teenager, man, young woman, ...)"
 Scene: "[spacious, rustic, urban, ...]" HOIPrompt: "Photo of a (black) (young woman) making and cooking a sandwich, (urban), (4k), (backlit), (partial view), (Canon Eos 5D)"
 Shooting: "(DSLR, grainy, 4k, ...), (warm lighting, blue hour, backlit, ...), (partial view, back view, ...), (Canon Eos 5D, iPhone 12, ...)"

(a) Examples of HOIPrompts



(b) Examples of generated images and annotations

Figure 2. Illustration of a) HOIPrompts and b) how HOIPrompts guide the text-to-image generation process to enhance diversity. For more visualization, please refer to the Appendix.

offers three key advantages: **(1) Rich synonymous descriptions:** Traditional datasets often use specific descriptions to annotate objects and relationships. In contrast, *Magic-HOI* provides a multitude of synonymous descriptions across datasets. For example, it distinguishes between the object ‘person’ and ‘woman’/‘man’, as well as the action descriptions ‘read’ and ‘look at’. This diversity enhances the model’s robustness and adaptability. **(2) Large-scale data:** Training large models necessitates diverse data. Single datasets, such as the 20K images in HICO-DET, can cause the model underfitting. *Magic-HOI* make a significant effort to build a large-scale dataset for effective training, containing an extensive over 186K images. **(3) Rich object/actions/interactions categories:** Current models are primarily trained on data with 80 objects and 117 verbs. However, *Magic-HOI* significantly broadens the scope by offering over 2.4K objects, 1.2K actions, and 20K HOI interactions. Such a large vocabulary greatly enhances the model’s generalization capabilities.

3.2. SynHOI

As shown in Fig. 1-(b), *Magic-HOI* exhibits an inherent long-tail issue. Thus, we present a high-quality synthetic HOI dataset called *SynHOI* to complement *Magic-HOI*. To make the flow of dataset production scalable, we present an automatic pipeline, including the HOIPrompt design, automatic labeling and filtering, and quality verification, designed to continually scale up the generation of diverse and high-precision HOI-annotated data, as shown in Fig. 2. Please refer to the Appendix for details.

SynHOI has three key data characteristics: **(1) High-quality data.** *SynHOI* showcases high-quality HOI annotations. First, we employ CLIPScore [16] to measure the similarity between the synthetic images and the corresponding HOI triplet prompts. The *SynHOI* dataset achieves a high CLIPScore of 0.849, indicating a faithful reflection of the HOI triplet information in the synthetic images. Second, Fig. 2-(b) provides evidence of the high quality of detection annotations in *SynHOI*, attributed to the effectiveness of the SOTA detector [64] and the alignment of *SynHOI* with real-

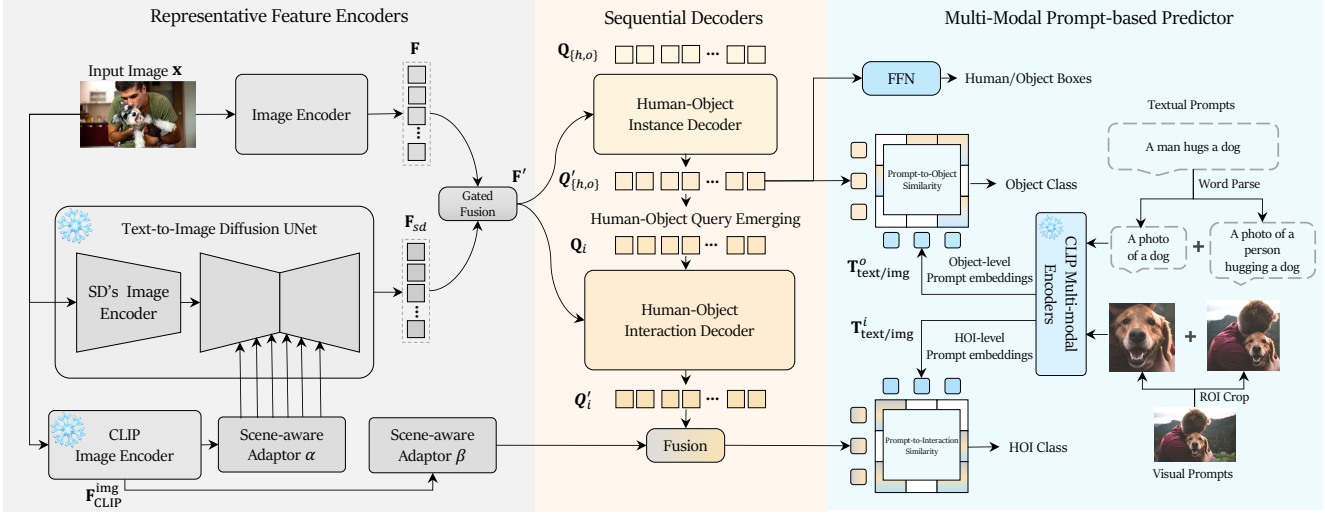


Figure 3. Overview of *MP-HOI*, comprising three components: Representative Feature Encoder, Sequential Instance and Interaction Decoders, and Multi-modal Prompt-based Predictor. Ultimately, it can leverage textual or visual prompts to detect open-world HOIs.

world data distributions. The visualization of *SynHOI* is presented in the Appendix. (2) **High-diversity data.** *SynHOI* exhibits high diversity, offering a wide range of visually distinct images. Fig. 2-(b) demonstrates the impact of random variations in person’s descriptions, environments, and photographic information within the HOIPrompts on the diversity of synthetic images. (3) **Large-scale data with rich categories.** *SynHOI* aligns *Magic-HOI*’s category definitions to effectively address the long-tail issue in *Magic-HOI*. It consists of over 100K images, 130K person bounding boxes, 140K object bounding boxes, and 240K HOI triplet instances.

4. Methodology

4.1. The Overview of MP-HOI

MP-HOI is a multi-modal prompt-based HOI detector comprising three components: Representative Feature Encoder, Sequential Instance and Interaction Decoders, and Multi-modal Prompt-based Predictor, as shown in Fig. 3.

Given an input image \mathbf{x} , the Representative Feature Encoder firstly encodes it to powerful feature representations \mathbf{F}' , which have fused the vanilla image features \mathbf{F} from the image encoder and the HOI-associated representations \mathbf{F}_{sd} obtained from text-to-image diffusion model [44] (refer to Sec. 4.2 for details). Based on \mathbf{F}' , the Sequential Instance and Interaction Decoders are designed to decode the pairwise human-object representations and their corresponding interaction representations in a sequential manner. Specifically, the instance decoder takes the input of initialized pairwise human-object queries $\mathbf{Q}_{\{h,o\}} \in \mathbb{R}^{2 \times N \times C}$ and produces the updated queries $\mathbf{Q}'_{\{h,o\}} \in \mathbb{R}^{2 \times N \times C}$, where N is the number of paired queries and C indicates the channel dimension. Subsequently, we merge the human-object

queries $\mathbf{Q}'_{\{h,o\}}$, into interaction queries $\mathbf{Q}_i \in \mathbb{R}^{N \times C}$ via average pooling. The interaction decoder then receives these interaction queries and updates them to $\mathbf{Q}'_i \in \mathbb{R}^{N \times C}$. To enhance the interaction representations, we integrate the global CLIP image representation into each of the interaction queries $\mathbf{Q}'_i \in \mathbb{R}^{N \times C}$ via element-wise addition.

The complete HOI prediction for an image includes human and object bounding boxes, object class, and interactions class. Firstly, we employ Feed-Forward Network (FFN) on $\mathbf{Q}'_{h,o}$ to predict the human and object bounding box. Secondly, we introduce the Multi-modal Prompt-based Predictor to identify the object and interaction classes (refer to Sec. 4.3).

4.2. Representative Feature Encoder

Motivation. As illustrated in Sec. 3, even though we have unified six datasets into *Magic-HOI* and introduced the synthetic dataset *SynHOI*, compared with the amount of training data for large-scale visual-linguistic models, such as CLIP [42] or Stable Diffusion (SD) [44], we still have 1,000 times less data. We firmly believe that such models can further provide effective encoded image representations for training instance/interaction decoders. Specifically, considering that the SD model utilizes cross-attention mechanisms between text embeddings and visual representations for text-conditioned image generation, we assume the SD model has a substantial correlation between their feature spaces and semantic concepts in language. To further demonstrate this, we employ DAAM [49] to generate pixel-level linguistic concept heatmaps based on the image features extracted from the SD model. As depicted in Fig. 4, beyond the noun concepts highlighted in prior studies [28, 46, 57], we find that the internal representation space of a frozen SD model is highly relevant to verb con-

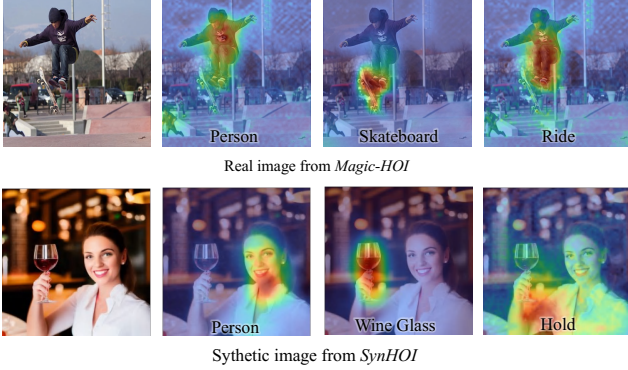


Figure 4. Attribution analysis of Stable Diffusion between object/interaction texts and real/synthetic images. The visualization solely utilizes time-step 0.

cepts and their associated contexts, which are very useful for the HOI task.

Extract Local Semantic Association in Stable Diffusion.

Inspired by the above observations, the feature maps derived from the SD model are expected to exhibit superior semantic correspondence within local regions, establishing an association between the semantic information of the text and the corresponding regions within the image. Accordingly, we utilize the UNet-based diffusion model to extract the feature maps, which are not only associated with the noun concepts but also include the verb concepts and the corresponding contextual details. Different from using a series of denoising steps to generate an image, we directly feed the input image x into the UNet and perform a single forward pass via the network. The output multi-scale image features are as $F_{sd} = \text{SD}(x, A_x)$, where A_x denotes the text-related representation that corresponds to x . Typically, A_x can be obtained by utilizing a text encoder, such as the CLIP text encoder, to encode the description of the image. However, as a discriminative task, the description of the HOI image is not available in the inference. To address this, we replace it with the global feature representation of x extracted from the CLIP image encoder, which bears significant relevance to the textual representation.

Two Scene-aware Adaptors. Given that the CLIP image features F_{clip}^{img} represents a global feature, it inherently encompasses the contextual information of the overall scene. As shown in Fig. 3, we introduce two scene-aware adaptors, denoted as α and β , to project F_{clip}^{img} into feature spaces more consistent with the SD model and Multi-modal Prompt-based Predictor. Regarding the scene-aware adaptor α , since the CLIP model is trained to align global visual and textual representations within a latent space, the F_{clip}^{img} can be employed as a substitute for the textual representation. Hence, we can train an MLP to directly project F_{clip}^{img} to a textual space suitable for the SD model. As for the scene-aware adaptor β , we project F_{clip}^{img} through it and incorporate the adapted version into each interaction query of Q'_i . This

adjustment allows for the tuning of these interaction queries to align more effectively with multi-modal prompts.

Fused Features for Better HOI Representing. Following the previous detectors [30, 61], we also employ a common image encoder to obtain the vanilla image features F . Then, we perform Gated Fusion to generate more informative F' . Specifically, F' combines F with supplementary features F_{sd} for the interaction decoder, while retaining the original features F for the instance decoder.

4.3. Multi-Modal Prompt-based Predictor

Motivation. As discussed in the introduction, existing language-guided-only HOI detectors face challenges in generalization when dealing with entirely new categories that lack any relevant text descriptions in the training data. They also struggle to address complex scenarios with high interaction ambiguity. Therefore, we are motivated to incorporate visual prompts to enhance the effectiveness and generalization capability in open-world scenarios. Visual prompts can act as exemplars, allowing the model to detect the same objects/interactions, and they also aid in reducing semantic ambiguity present in textual prompts. Simultaneously, textual prompts provide semantic context to enhance the understanding of visual prompts. Given that CLIP [42] is trained to align visual and textual representations through hundreds of millions of image-text pairs, we naturally leverage its pre-trained image and text encoders to encode both textual and visual prompts.

Encoding Multi-Modal Prompts. For textual prompt inputs, we can parse them into instance descriptions like "A photo of a [object]," and interaction descriptions like "A photo of a [man/child/...] [verb-ing] a [object]." These sentences can then be encoded to obtain object category embeddings T_{text}^o or interaction category embeddings T_{text}^i . For visual prompt inputs, we can crop them into corresponding object ROI images and HOI ROI images using user-provided bounding boxes. These images can also be used to obtain object category embeddings T_{img}^o or interaction category embeddings T_{img}^i .

Multi-Modal Prompt-based Classifiers. Formally, the final object category distributions P_o and HOI category distributions P_i can be calculated as,

$$P_o = \text{softmax}(Q'_o * T_{text/img}^o{}^T) \quad (1)$$

$$P_i = \text{softmax}(Q'_i * T_{text/img}^i{}^T) \quad (2)$$

where $Q'_o \in \mathbb{R}^{N \times C}$ and $Q'_i \in \mathbb{R}^{N \times C}$ denotes object queries and interaction queries, respectively, and softmax indicates the row-wise softmax operation.

4.3.1 Reformulate the Optimization of HOI

Cross-modal Contrastive Learning. *MP-HOI* reformulate the HOI task as a similarity learning process between multi-modal prompts and objects/interactions via a unified contrastive loss. Thus, it can perform multi-modal prompt-to-object/interaction alignments for large-scale image-prompt pairs in *Magic-HOI* and *SynHOI*. In general, given B image-prompt pairs $\{(\mathbf{x}_m, \text{prompt}_m)\}_{m=1}^B$ in our datasets, the unified contrastive loss \mathcal{L}_c can be formulated as,

$$\mathcal{L}_c = -\frac{1}{B} \sum_{m=1}^B \log\left(\frac{\exp(S(\mathbf{x}_m, \text{prompt}_m)/\tau)}{\sum_{n=1}^B \exp(S(\mathbf{x}_m, \text{prompt}_n)/\tau)}\right) \quad (3)$$

where $S(\mathbf{x}_m, \text{prompt}_n)$ is prompt-to-image similarity between m -th image \mathbf{x}_m and n -th prompt prompt_n . τ is a temperature to scale the logits. Through \mathcal{L}_c , we could obtain the object contrastive loss \mathcal{L}_c^o and the interaction contrastive loss \mathcal{L}_c^i , which enable the model to learn generalizable and transferable objects/interactions representations from the large-scale categories.

Overall Loss. Following the query-based methods [30, 61], we employ the Hungarian algorithm to match predictions to each ground-truth. The overall loss is computed between the matched predictions and their corresponding ground-truths, which includes the box regression loss \mathcal{L}_b , the intersection-over-union loss \mathcal{L}_g , the object contrastive loss \mathcal{L}_c^o , and the interaction contrastive loss \mathcal{L}_c^i ,

$$\mathcal{L} = \lambda_b \mathcal{L}_b + \lambda_g \mathcal{L}_g + \lambda_c^o \mathcal{L}_c^o + \lambda_c^i \mathcal{L}_c^i, \quad (4)$$

where \mathcal{L}_b and \mathcal{L}_g contain both human and object localization. λ_b , λ_g , λ_c^o and λ_c^i are used to adjust the weights of each loss component.

4.4. Training and Inference Details

Training Details. With the unified cross-modality contrastive loss that we have defined, we are able to jointly optimize multi-modal prompts. Considering the training cost, we employ a 50% probability to randomly select either a visual prompt or a textual prompt for each iteration. To use visual prompts for training, we sample two images containing a person paired with the exact same single/multiple objects and single/multiple interaction categories from the dataset, and then choose one of them as the visual prompt.

Inference Details. Users can freely choose from available prompts. We present quantitative results based on textual prompts to ensure a fair comparison.

5. Experiments

5.1. Experimental Settings

Datasets and Evaluation Metrics. We evaluate our models on four benchmarks: HICO-DET [6], V-COCO [13],

Method	Backbone	Default		
		Full	Rare	Non-Rare
<i>Expert Models</i>				
QPIC [48]	ResNet-50	29.07	21.85	31.23
CDN [61]	ResNet-50	31.44	27.39	32.64
DOQ [41]	ResNet-50	33.28	29.19	34.50
IF [32]	ResNet-50	33.51	30.30	34.46
GEN-VLKT [30]	ResNet-50	33.75	29.25	35.10
PViC [63]	ResNet-50	<u>33.80</u>	<u>29.28</u>	<u>35.15</u>
QAHOI [8]	Swin-L	35.78	29.80	37.56
FGAHOI [35]	Swin-L	37.18	30.71	39.11
PViC [63]	Swin-L	<u>43.35</u>	<u>42.25</u>	<u>43.69</u>
<i>Pre-trained Models</i>				
RLIPv1 [59]	ResNet-50	13.92	11.20	14.73
RLIPv1 [59] [†]	ResNet-50	30.70	24.67	32.50
RLIPv2 [60]	ResNet-50	17.79	19.64	17.24
RLIPv2 [60] [†]	ResNet-50	35.38	29.61	37.10
<i>Generalist Models</i>				
<i>MP-HOI-S</i>	ResNet-50	36.50 ^{↑2.70}	35.48 ^{↑6.20}	36.80 ^{↑1.65}
<i>MP-HOI-L</i>	Swin-L	44.53^{↑1.18}	44.48^{↑2.23}	44.55^{↑0.86}

¹ RLIP v2: VG+COCO+O365, 2200K+ pretraining data.

² *MP-HOI*: *Magic-HOI*+*SynHOI*, 286K+ training data.

Table 2. Performance comparison on HICO-DET in terms of mAP (%). † indicates the models are fully fine-tuned on HICO-DET. The results highlighted in underlined represent state-of-the-art performance among expert models, which we primarily compare to.

Method	Backbone	Default		
		Full	Rare	Non-Rare
<i>Expert Models</i>				
GEN-VLKT [30]	ResNet-50	10.87	10.41	20.91
<i>Generalist Models</i>				
<i>MP-HOI-S</i>	ResNet-50	12.61	14.78	20.28
<i>MP-HOI-L</i>	Swin-L	16.21	18.59	25.76

Table 3. Performance comparison on SWiG in terms of mAP (%). Given that GEN-VLKT is one of the most powerful and representative expert models, we train it on SWiG for comparison.

SWiG [52] and HCVRD [68]. The mean Average Precision (mAP) is used as the evaluation metric, following standard protocols [30, 61].

Implementation Details. We implement two variant architectures of *MP-HOI*: *MP-HOI-S*, and *MP-HOI-L*, where ‘S’ and ‘L’ refer to small and large, respectively. For *MP-HOI-S*, we use ResNet-50 as the backbone and a six-layer vanilla Transformer encoder [5] as the feature extractor. Both the human-object decoder and interaction decoder are three-layer vanilla Transformer decoders. For *MP-HOI-L*, we employ Swin-L [34] as the backbone, with Transformer decoders of six layers. The CLIP-based model and diffusion model are frozen during training for all the settings. For a fair comparison with previous studies, we employ the ViT-B variant of the CLIP model within the *MP-HOI-S*. Conversely, to enhance performance in real-world scenarios, the ViT-L variant is utilized within the *MP-HOI-L*.

5.2. General HOI Detection

MP-HOI is a generalist HOI detector trained on the unified dataset-*Magic-HOI*. Since *Magic-HOI* unifies multi-

Method	Backbone	Default		
		Full	Rare	Non-Rare
<i>Expert Models</i>				
GEN-VLKT [30]	ResNet-50	6.58	5.81	14.02
<i>Generalist Models</i>				
MP-HOI-S	ResNet-50	8.08	6.86	14.31
MP-HOI-L	Swin-L	11.29	9.01	18.68

Table 4. Performance comparison on HCVRD in terms of mAP (%). We train GEN-VLKT on HCVRD for comparison.

Method	Backbone	AP (Scenario 1)	AP (Scenario 2)
<i>Expert Models</i>			
IDN [27]	ResNet-50	53.3	60.3
FCL [19]	ResNet-50	52.4	-
SCG [62]	ResNet-50	54.2	60.9
HOTR [21]	ResNet-50	55.2	64.4
QPIC [48]	ResNet-50	58.8	61.0
CDN [61]	ResNet-50	61.7	63.8
GEN-VLKT [30]	ResNet-50	62.4	64.5
<i>Generalist Models</i>			
<i>Zero-shot Testing</i>			
MP-HOI-S	ResNet-50	37.5	44.2
<i>10% Training Data</i>			
MP-HOI-S	ResNet-50	57.7	60.2
<i>100% Training Data</i>			
MP-HOI-S	ResNet-50	66.2	67.6

Table 5. Performance comparison on V-COCO.

ple datasets, we could directly evaluate *MP-HOI* on these benchmarks. We compare its performance with two types of models: expert models and pretrained models. As there are no standard results available like mAP on SWiG-HOI and HCVRD datasets for comparison, we train the GEN-VLKT [30], which is one of the representative expert models, as the baseline for our comparison.

HICO-DET. Compared with the expert models, *MP-HOI* outperforms them by a significant margin under both ResNet-50 and Swin-L backbones, notably achieving a remarkable 6.20 improvement in mAP in the rare setting. Compared with the pre-trained model, *MP-HOI* also surpasses RLIPv2 [60], which is pretrained on over 2200K data and fine-tuned on HICO-DET. Finally, *MP-HOI-L* achieves a new SOTA performance with 44.53 mAP.

SWiG-HOI & HCVRD. Compared with GEN-VLKT [30], *MP-HOI* exhibits superior performance on these two datasets that contain more HOI concepts. This highlights *MP-HOI*'s capability to serve as a generalist HOI detector.

5.3. Zero-Shot HOI Detection

We demonstrate the generalization of the model trained on our proposed *Magic-HOI* and *SynHOI* datasets by evaluating it on the unseen dataset-V-COCO. As shown in Tab. 5, *MP-HOI* exhibits excellent zero-shot performance on V-COCO. After fine-tuning the model on just 10% of the training data, it achieves an impressive AP of 57.7, surpassing some expert models that are fully trained on V-COCO. Moreover, when utilizing the full training data, *MP-HOI*

outperforms all expert models by a significant margin.

5.4. Ablation Study

Two Key Representations. As described in Sec. 4.1, we introduce \mathbf{F}_{sd} from stable diffusion and \mathbf{F}_{clip}^{img} from CLIP image encoder, into *MP-HOI*. Tab. 6 demonstrates their effectiveness. Initially, we establish a baseline model following the design of GEN-VLKT [30] while excluding the knowledge distillation component. Subsequently, the incorporation of \mathbf{F}_{sd} leads to significant improvements, particularly in terms of rare mAP. This suggests that the internal features in stable diffusion play a crucial role in enhancing the representation of rare categories in HICO-DET. Additionally, the integration of \mathbf{F}_{clip}^{img} further improves the non-rare AP, as it aligns the queries with multi-modal prompts.

Diffusion Time Steps. We investigate the effectiveness of different diffusion steps in extracting \mathbf{F}_{sd} , similarly to [3, 57]. Diffusion models control the noise distortion added to the input image by varying the value of t . Stable diffusion [44] uses a total of 1000 time steps. We set t values to 0, 100, 500 for ablation studies. As demonstrated in Tab. 7, the best performance is achieved when $t = 0$. It is worth noting that using \mathbf{F}_{sd} extracted from input images with higher noise levels would decrease performance and potentially impact the learning of interactions.

Two Scene-Aware Adaptors. Tab. 8 illustrates that both scene-aware adaptors effectively facilitate Stable Diffusion and the CLIP image encoder in extracting feature representations aligned with our *MP-HOI* framework and HOI task, thereby improving performance.

Contrastive Learning. As in Sec. 4.3.1, we introduce the object contrastive loss \mathcal{L}_c^o and the interaction contrastive loss \mathcal{L}_c^i . We leverage these two losses to facilitate alignment between multi-modal prompts and objects/interactions. The results in Tab. 9 demonstrate the effectiveness of both loss items, particularly in improving performance on rare classes, where the model exhibits improved discrimination between objects and interactions.

Multi-modal Prompts. As shown in Tab. 10, integrating visual prompts during training could enhance performance compared to using only textual prompts. This improvement is due to the inherent semantic ambiguity in textual prompts, which visual prompts help to reduce.

The Scale of Training Data. Beyond solely training on HICO-DET, we incorporate our *Magic-HOI* dataset that includes five additional datasets for training. As in Tab. 11, this leads to improved performance and expands the range of object and interaction categories covered. Furthermore, the inclusion of our *SynHOI* dataset could also effectively enhance performance, particularly on rare classes.

F_{sd}	F_{clip}^{img}	Full	Rare	Non-rare
		31.99	29.63	32.70
✓		32.92	31.29	33.41
✓	✓	34.41	31.07	35.40

Table 6. Ablation results of two key representations on HICO-DET.

\mathcal{L}_c^o	\mathcal{L}_c^i	Full	Rare	Non-Rare
		34.41	31.07	35.40
✓		34.56	31.29	35.43
✓	✓	34.82	31.87	35.49

Table 9. The ablation study of contrastive losses at two levels on HICO-DET.

Time step	Full	Rare	Non-Rare
0	34.41	31.07	35.40
100	34.03	30.58	35.02
500	33.59	29.80	34.71

Table 7. Ablation results of different diffusion time steps on HICO-DET.

T.	V.	Full	Rare	Non-Rare
✓		34.82	31.87	35.49
✓	✓	35.18	32.16	35.57

Table 10. The ablation study of multi-modal prompts on HICO-DET. T. and V. denotes textual prompt and visual prompt, respectively.

Adaptor α	Adaptor β	Full	Rare	Non-Rare
✓		34.08	30.64	35.08
	✓	33.76	29.89	34.82
✓	✓	34.41	31.07	35.40

Table 8. Ablation results of two scene adaptors on HICO-DET.

Training Data	Full	Rare	Non-Rare
HICO-DET	35.18	32.16	35.57
<i>Magic-HOI</i>	35.93	34.68	36.26
+ <i>SynHOI</i>	36.50	35.48	36.80

Table 11. The impact of the scale of training data.




Figure 5. In-the-wild test based on arbitrary textual prompts. Each HOI triplet is represented in the same color.



Figure 6. In-the-wild test based on arbitrary visual prompts. The first row is the visual prompt (green boxes define objects and red boxes define interactions), and the second row is the corresponding detection results. We highlight the powerful HOI detection performance, the capability to detect single HOIs (a, b, c), as well as a specific HOI composition (d), and the flexibility to integrate the visual prompt for interaction definition and textual prompt for object definitions (e.g., woman and phone) for the test (e).

5.5. HOI Detection in the Open World

In addition to the zero-shot experiments, we also conduct open-world testing on arbitrary images from the Internet as shown in Fig. 5 and Fig. 6. Our *MP-HOI* could leverage arbitrary textual and visual prompts to detect HOIs. Interestingly, we find *MP-HOI* could support extensive textual prompts, such as Princess Diana and Prince Charles for human description in Fig. 5-(a) and Döner Kebab for object description in Fig. 5-(d). Also, as in Fig. 6, *MP-HOI* could detect a specific HOI composition via visual prompts, and has the flexibility to integrate the visual prompt for interaction definition and textual prompt for object definitions (e.g., woman and phone) for the test as in Fig. 6-(e).

6. Conclusion

We introduce *MP-HOI*, a powerful multi-modal prompt-based HOI detector. It leverages both textual descriptions and visual exemplars to realize HOI detection in the open world. We introduce a large-scale unified HOI

dataset named *Magic-HOI*, and a high-quality synthetic HOI dataset called *SynHOI* for effective training. Our results demonstrate that the trained *MP-HOI*, as a generalist HOI detector, exhibits remarkable zero-shot capability in real-world scenarios and consistently achieves a new state-of-the-art performance across various benchmarks. We hope our models and datasets can serve the community for further research.

Acknowledgment

The work is partially supported by the Young Scientists Fund of the National Natural Science Foundation of China under grant No.62106154, by the Natural Science Foundation of Guangdong Province, China (General Program) under grant No.2022A1515011524, and by Shenzhen Science and Technology Program JCYJ20220818103001002 and ZDSYS20211021111415025, and by the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong (Shenzhen).

References

- [1] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. [2](#)
- [2] 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](#), [3](#)
- [3] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khruikov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021. [7](#)
- [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. [3](#)
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*, pages 213–229. Springer, 2020. [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] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 2018. [1](#), [3](#)
- [8] Junwen Chen and Keiji Yanai. Qahoi: query-based anchors for human-object interaction detection. *arXiv preprint arXiv:2112.08647*, 2021. [3](#), [6](#)
- [9] 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. [3](#)
- [10] Hao-Shu Fang, Yichen Xie, Dian Shao, and Cewu Lu. Dirv: Dense interaction region voting for end-to-end human-object interaction detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1291–1299, 2021. [2](#)
- [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 *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 696–712. Springer, 2020. [2](#)
- [13] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015. [6](#)
- [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](#), [3](#)
- [15] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *2015 IEEE conference on computer vision and pattern recognition (CVPR)*, pages 961–970. IEEE, 2015. [1](#)
- [16] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. [3](#)
- [17] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pages 584–600. Springer, 2020. [1](#), [3](#)
- [18] 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.
- [19] 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](#), [3](#), [7](#)
- [20] Bumsoo Kim, Taeho Choi, Jaewoo Kang, and Hyunwoo J Kim. Uniondet: Union-level detector towards real-time human-object interaction detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pages 498–514. Springer, 2020. [2](#)
- [21] 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. [3](#), [7](#)
- [22] 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*, pages 2925–2934, 2023. [3](#)
- [23] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. [3](#)
- [24] Yong-Lu Li, Liang Xu, Xinpeng Liu, Xijie Huang, Yue Xu, Mingyang Chen, Ze Ma, Shiyi Wang, Hao-Shu Fang, and Cewu Lu. Hake: Human activity knowledge engine. *arXiv preprint arXiv:1904.06539*, 2019. [3](#)
- [25] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactiveness knowledge for human-object interaction

- detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019. [2](#)
- [26] Yong-Lu Li, Xinpeng Liu, Han Lu, Shiyi Wang, Junqi Liu, Jiefeng Li, and Cewu Lu. Detailed 2d-3d joint representation for human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10166–10175, 2020. [2](#)
- [27] Yong-Lu Li, Xinpeng Liu, Xiaoqian Wu, Yizhuo Li, and Cewu Lu. Hoi analysis: Integrating and decomposing human-object interaction. *Advances in Neural Information Processing Systems*, 33:5011–5022, 2020. [7](#)
- [28] Ziyi Li, Qinye Zhou, Xiaoyun Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Guiding text-to-image diffusion model towards grounded generation. *arXiv preprint arXiv:2301.05221*, 2023. [4](#)
- [29] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jia-shi 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](#)
- [30] Yue Liao, Aixi Zhang, Miao Lu, Yongliang Wang, Xiaobo Li, and Si Liu. Gen-vlkt: Simplify association and enhance interaction understanding for hoi detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20123–20132, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [31] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection—a new baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6536–6545, 2018. [1](#)
- [32] 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. [6](#)
- [33] Ye Liu, Junsong Yuan, and Chang Wen Chen. Consnet: Learning consistency graph for zero-shot human-object interaction detection. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4235–4243, 2020. [2](#), [3](#)
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. [6](#)
- [35] Shuailei Ma, Yuefeng Wang, Shanze Wang, and Ying Wei. Fgahoi: Fine-grained anchors for human-object interaction detection. *arXiv preprint arXiv:2301.04019*, 2023. [3](#), [6](#)
- [36] Shan Ning, Longtian Qiu, Yongfei Liu, and Xuming He. Hoiclip: Efficient knowledge transfer for hoi detection with vision-language models. *arXiv preprint arXiv:2303.15786*, 2023. [1](#), [3](#)
- [37] Jeeseung Park, Jin-Woo Park, and Jong-Seok Lee. Viplo: Vision transformer based pose-conditioned self-loop graph for human-object interaction detection. *arXiv preprint arXiv:2304.08114*, 2023. [2](#)
- [38] Julia Peyre, Ivan Laptev, Cordelia Schmid, and Josef Sivic. Detecting unseen visual relations using analogies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1981–1990, 2019. [3](#)
- [39] PicDataset. Picdataset.com: Artificial intelligence news today, 2023. [3](#)
- [40] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 401–417, 2018. [2](#)
- [41] Xian Qu, Changxing Ding, Xingao Li, Xubin Zhong, and Dacheng Tao. Distillation using oracle queries for transformer-based human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19558–19567, 2022. [6](#)
- [42] 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. [1](#), [3](#), [4](#), [5](#)
- [43] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [2](#)
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [4](#), [7](#)
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. [2](#)
- [46] Jordan Shipard, Arnold Wiliem, Kien Nguyen Thanh, Wei Xiang, and Clinton Fookes. Diversity is definitely needed: Improving model-agnostic zero-shot classification via stable diffusion. [4](#)
- [47] 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 (CVPR)*, pages 10410–10419, 2021. [3](#)
- [48] 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. [3](#), [6](#), [7](#)
- [49] Raphael Tang, Akshat Pandey, Zhiying Jiang, Gefei Yang, Karun Kumar, Jimmy Lin, and Ferhan Ture. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*, 2022. [4](#)
- [50] Bo Wan, Desen Zhou, Yongfei Liu, Rongjie Li, and Xuming He. Pose-aware multi-level feature network for human object interaction detection. In *Proceedings of the IEEE/CVF*

- International Conference on Computer Vision*, pages 9469–9478, 2019. 1, 2
- [51] Hai Wang, Wei-shi Zheng, and Ling Yingbiao. Contextual heterogeneous graph network for human-object interaction detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, pages 248–264. Springer, 2020. 2
- [52] 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
- [53] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. Learning human-object interaction detection using interaction points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4116–4125, 2020. 2
- [54] Mingrui Wu, Jiabin 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. 1, 3
- [55] Xiaoqian Wu, Yong-Lu Li, Xinpeng Liu, Junyi Zhang, Yuzhe Wu, and Cewu Lu. Mining cross-person cues for body-part interactiveness learning in hoi detection. In *Proceedings of the European Conference on Computer Vision*, pages 121–136. Springer, 2022. 1
- [56] Bingjie Xu, Yongkang Wong, Junnan Li, Qi Zhao, and Mohan S Kankanhalli. Learning to detect human-object interactions with knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [57] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. *arXiv preprint arXiv:2303.04803*, 2023. 4, 7
- [58] Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. Physical inertial poser (pip): Physics-aware real-time human motion tracking from sparse inertial sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13167–13178, 2022. 1
- [59] Hangjie Yuan, Jianwen Jiang, Samuel Albanie, Tao Feng, Ziyuan Huang, Dong Ni, and Mingqian Tang. Rlip: Relational language-image pre-training for human-object interaction detection. *Advances in Neural Information Processing Systems*, 35:37416–37431, 2022. 3, 6
- [60] Hangjie Yuan, Shiwei Zhang, Xiang Wang, Samuel Albanie, Yining Pan, Tao Feng, Jianwen Jiang, Dong Ni, Yingya Zhang, and Deli Zhao. Rlipv2: Fast scaling of relational language-image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21649–21661, 2023. 6, 7
- [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, 3, 5, 6, 7
- [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*, pages 13319–13327, 2021. 2, 7
- [63] Frederic Z Zhang, Yuhui Yuan, Dylan Campbell, Zhuoyao Zhong, and Stephen Gould. Exploring predicate visual context in detecting of human-object interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10411–10421, 2023. 6
- [64] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 3
- [65] Long Zhao, Liangzhe Yuan, Boqing Gong, Yin Cui, Florian Schroff, Ming-Hsuan Yang, Hartwig Adam, and Ting Liu. Unified visual relationship detection with vision and language models. *arXiv preprint arXiv:2303.08998*, 2023. 3
- [66] Desen Zhou, Zhichao Liu, Jian Wang, Leshan Wang, Tao Hu, Errui Ding, and Jingdong Wang. Human-object interaction detection via disentangled transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19568–19577, 2022. 1, 3
- [67] Yufan Zhou, Ruiyi Zhang, Changyou Chen, Chunyuan Li, Chris Tensmeyer, Tong Yu, Jiuxiang Gu, Jinhui Xu, and Tong Sun. Towards language-free training for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17907–17917, 2022. 2
- [68] Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, and Anton van den Hengel. Hcvrd: A benchmark for large-scale human-centered visual relationship detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. 3, 6