
Visual Diversity and Region-aware Prompt Learning for Zero-shot HOI Detection

Chanyeong Yang¹ Taejoon Song² Jihwan Park¹ Hyunwoo J. Kim^{2*}

¹Korea University ²Korea Advanced Institute of Science and Technology

{0814gerrardso, jseven7071}@korea.ac.kr {taehoont, hyunwoojkim}@kaist.ac.kr

Abstract

Zero-shot Human-Object Interaction detection aims to localize humans and objects in an image and recognize their interaction, even when specific verb-object pairs are unseen during training. Recent works have shown promising results using prompt learning with pretrained vision-language models such as CLIP, which align natural language prompts with visual features in a shared embedding space. However, existing approaches still fail to handle the *visual complexity of interaction*—including (1) *intra-class visual diversity*, where instances of the same verb appear in diverse poses and contexts, and (2) *inter-class visual entanglement*, where distinct verbs yield visually similar patterns. To address these challenges, we propose **VDRP**, a framework for *Visual Diversity and Region-aware Prompt learning*. First, we introduce a visual diversity-aware prompt learning strategy that injects group-wise visual variance into the context embedding. We further apply Gaussian perturbation to encourage the prompts to capture diverse visual variations of a verb. Second, we retrieve region-specific concepts from the human, object, and union regions. These are used to augment the diversity-aware prompt embeddings, yielding region-aware prompts that enhance verb-level discrimination. Experiments on the HICO-DET benchmark demonstrate that our method achieves state-of-the-art performance under four zero-shot evaluation settings, effectively addressing both intra-class diversity and inter-class visual entanglement. Code is available at <https://github.com/mlvlab/VDRP>.

1 Introduction

Human-Object Interaction (HOI) detection [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13] aims to localize humans and objects in an image and recognize the interactions between them, serving as a cornerstone for fine-grained scene understanding. Unlike standard HOI detection, which assumes supervision over all interactions, zero-shot HOI detection must generalize to unseen combinations of verbs and objects. While recent advances in Vision-Language Models (VLMs) have significantly improved the zero-shot recognition of object and attribute classes [14, 15, 16, 17, 18, 19, 20], zero-shot HOI detection remains fundamentally more challenging. As previous works [21, 22] pointed out, the difficulty stems not only from the compositional novelty of interactions but also from the visual complexity of interactions—where each verb class exhibits large intra-class visual diversity, and different verbs frequently produce visually similar patterns.

First, verb classes exhibit substantial intra-class diversity. As shown in Fig. 1-(A), instances of the verb “holding a baseball glove” may appear in drastically different poses, scales, or scene contexts, yet must be classified under the same label. To quantify this, we compute a diversity score using CLS features extracted from a frozen CLIP visual encoder. For verbs, we crop the union region; for objects, the object bounding box is used. The average pairwise cosine similarity within each class is

*Corresponding author

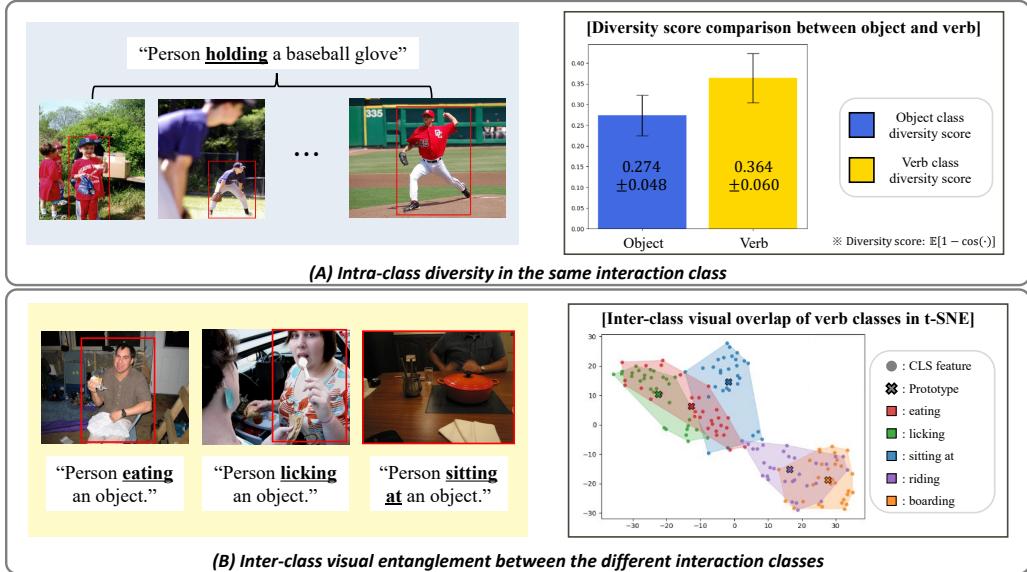


Figure 1: **Analysis of the visual complexity in HOI detection.** (A) Verb classes exhibit significant *intra-class visual diversity*, where instances of the same verb (e.g., “holding a baseball glove”) appear under varied poses, viewpoints, and scene contexts. To quantify this, we crop the union region and extract the CLIP visual CLS feature. A diversity score is then computed as the expected cosine dissimilarity $\mathbb{E}[1 - \cos(\cdot)]$ across samples of the same class. Verb classes exhibit higher diversity (0.364 ± 0.060) than object classes (0.274 ± 0.048), highlighting the difficulty of representing verbs with a single static embedding. (B) Verb classification also suffers from *inter-class visual entanglement*, where semantically distinct verbs (e.g., “eating”, “licking”, “sitting at”) yield visually similar patterns. To visualize this, we randomly select five verb classes, extract their union-region CLS features, and project them to 2D using t-SNE. The resulting clusters show significant overlap, highlighting the need for region-aware prompts to improve verb separability in HOI detection.

measured, and the diversity score is defined as the expectation of $1 - \cos(\cdot)$. This analysis shows that verbs have significantly higher intra-class diversity (0.364 ± 0.060) compared to objects (0.274 ± 0.048), indicating that a single prompt embedding is likely insufficient to capture such variation.

Second, interactions often exhibit *inter-class visual entanglement*, where different verbs produce highly similar visual patterns. As illustrated in Fig. 1-(B), semantically distinct verbs such as “eating”, “licking”, and “sitting at” frequently share similar human-object layouts. In such cases, accurate classification depends on regional differences—often localized in the human or object region—which global or union-level features may fail to capture. To analyze this phenomenon, we extract union-region CLS features from a frozen CLIP visual encoder across five verbs and project them into a 2D space using t-SNE. The resulting convex hulls show substantial overlap not only at the sample level but also across verb prototypes, revealing poor inter-class separability. This entanglement poses a major challenge for zero-shot HOI detection, where explicit class supervision is unavailable.

Recently, studies on zero-shot HOI detection have explored prompting strategies that leverage pretrained CLIP models [23, 24, 25, 26, 27]. These methods map HOI triplets to textual descriptions and embed them in a shared vision-language space. While effective for semantic alignment, most approaches assume a single static prompt per verb [23, 10, 26], making them inadequate for modeling the visual diversity within each class. Some incorporate spatial cues in the visual branch [24], but leave the text prompts agnostic to region-specific semantics. Others rely on LLM-generated descriptions [25], focusing on semantic differences across verbs but overlooking intra-class variation.

To overcome these limitations, we propose a new framework called **VDRP** (Visual Diversity and Region-aware Prompt learning) for zero-shot HOI detection. Our method enhances prompt representations in two complementary ways. First, we inject group-wise visual variance into the learnable context embeddings and apply Gaussian perturbation, allowing the prompts to reflect intra-class appearance diversity and better adapt to varied visual realizations during training. Second, we retrieve

region-specific concepts from the human, object, and union regions and use them to augment the prompt embeddings, producing region-aware prompts that improve verb discriminability.

Our **contributions** are summarized as follows:

- We propose a *visual diversity-aware prompt learning* that models intra-class variation by injecting group-wise variance into the context embedding and applying Gaussian perturbation, enabling the prompts to generalize across diverse appearances of the same verb.
- We introduce a *region-aware prompt augmentation* that leverages region-specific concept retrieval from human, object, and union regions to enrich the prompts, enhancing discriminability among visually similar verbs.
- We integrate the two modules into a unified framework, **VDRP**, and demonstrate its effectiveness on the HICO-DET benchmark, achieving new state-of-the-art performance under multiple zero-shot evaluation settings.

2 Related works

HOI detection. Human-Object Interaction (HOI) detection typically involves three sub-tasks: object detection, human-object pairing, and interaction classification. Thanks to advances in large-scale benchmarks [28, 29, 4, 21] and transformer-based architectures [30, 9, 31], a wide range of approaches have been proposed. Broadly, HOI methods fall into one-stage [23, 27, 4, 9, 11, 7] and two-stage [10, 26, 24, 25, 5, 6] paradigms. One-stage methods jointly predict object locations and interactions, often using set prediction frameworks such as DETR [30]. In contrast, two-stage methods decouple the task into object detection and interaction classification: a pre-trained detector localizes humans and objects, and a dedicated module predicts the verb label for each human-object pair. The division of the HOI task in two-stage approaches allows efficient training [5, 6], and shows promising results. Our work falls in the two-stage paradigm.

Zero-shot HOI detection. Zero-shot HOI detection aims to identify HOI triplets unseen during training, a task challenged by the long-tail distribution of compositional datasets. With the rise of Vision-Language Models (VLMs) [15, 32, 33, 14, 34] pretrained on large-scale image-text pairs, many works leverage their generalization for HOI. Several methods [10, 23, 27, 35] align HOI models with CLIP’s pretrained representations to enable effective transfer. Recent approaches [24, 26, 25] adopt prompt learning to adapt CLIP with few learnable parameters for fine-grained interaction understanding. However, most rely on a single static prompt per verb [10, 23, 26], limiting their ability to capture intra-class visual diversity. Later works add spatial cues or LLM-generated descriptions [24, 25], but still lack region-level adaptation or concept-level grounding. We address these gaps with visual diversity-aware prompt learning and region-aware prompt augmentation.

Prompt learning. Prompt learning is a widely used strategy for adapting vision-language models (VLMs) like CLIP to downstream tasks [36, 37, 38, 39, 40, 41, 42]. Early works such as CoOp [18] and CoCoOp [19] optimize learnable or image-conditioned context vectors, while MaPLe [38] jointly tunes visual and textual prompts to improve cross-modal alignment. In zero-shot HOI detection, CMMMP [24] and EZ-HOI [25] adopt multi-modal prompt learning, but do not account for visual diversity of verbs and use the same verb prompt across regions. This limits their ability to adapt to diverse verb appearances and handle visually similar interactions. Meanwhile, distribution-based prompt learning [37, 36, 40, 43] has shown that leveraging feature variance improves generalization. Inspired by this, we propose a method that injects group-wise visual variance—via modulation and perturbation—into prompt embeddings, and augments them with region-specific concepts, enabling effective handling of visual diversity and improving verb discriminability in zero-shot HOI detection.

3 Methods

In this section, we present our framework, **VDRP**, designed to address two key challenges in zero-shot Human-Object Interaction (HOI) detection: (1) *intra-class visual diversity*, where instances of the same verb exhibit a wide range of visual appearances, and (2) *inter-class visual entanglement*, where semantically distinct verbs appear visually similar. To tackle these challenges, VDRP introduces two complementary components. The first, *visual diversity-aware prompt learning*, addresses intra-class diversity by injecting group-wise visual variance into the context embeddings and applying variance-guided perturbation, resulting in *visual diversity-aware prompts* that better capture verb-specific

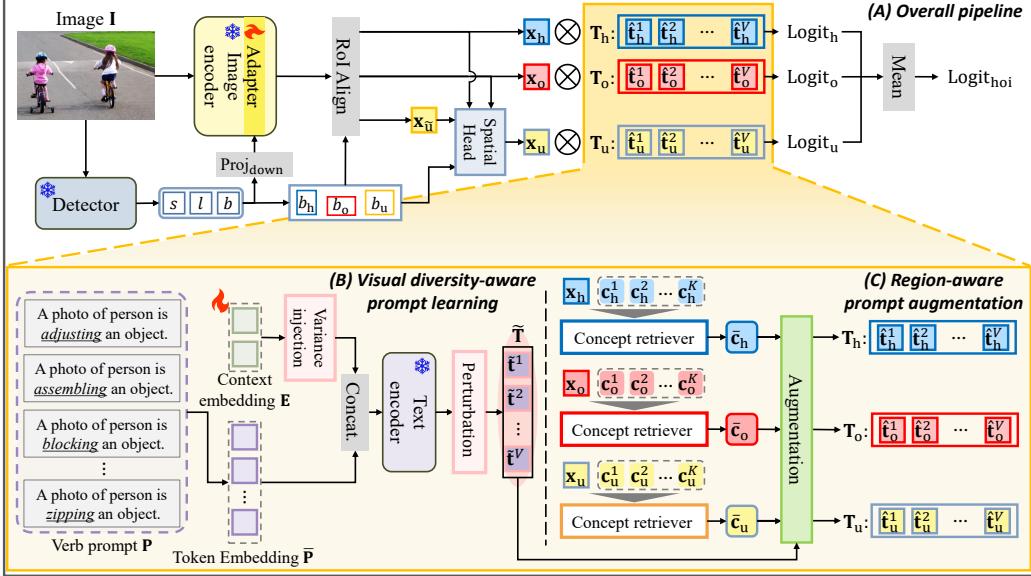


Figure 2: **Overview of our VDRP framework.** (A) We adopt a two-stage HOI detection pipeline with a frozen detector and a CLIP image encoder to extract human (x_h), object (x_o), and union ($x_{\bar{u}}$) features. A spatial head further refines the union feature into x_u for region-aware prompts via spatial encoding. (B) Visual diversity-aware prompts are generated by injecting group-wise variance and perturbation to model intra-class variation. (C) Retrieved region concepts are then fused with these prompts to produce final region-aware prompts T_h , T_o , and T_u used for verb classification.

appearance variations. The second, *region-aware prompt augmentation*, enhances these prompts using region-specific concepts retrieved from the human, object, and union regions, yielding *region prompts* that improve inter-class discriminability. We begin by outlining the overall pipeline (Section 3.1), followed by detailed descriptions of the two modules in Sections 3.2 and 3.3.

3.1 Overall pipeline

As illustrated in Fig. 2-(A), our method follows a two-stage HOI detection framework [10, 24, 5, 6], consisting of (1) human-object detection and (2) interaction classification. In the first stage, a frozen object detector (DETR [30]) is applied to the input image I to identify object instances. Each detection yields a triplet (s_n, l_n, b_n) , representing the confidence score, object class embedding, and bounding box of the n -th instance. In the second stage, we encode each instance into a prior embedding $\mathbf{p} = \text{Proj}_{\text{down}}([s; l; b]) \in \mathbb{R}^{d_{\text{down}}}$, which serves as guidance for task adaptation. Using this prior, we extract patch embeddings from a frozen CLIP encoder augmented with lightweight adapter layers inserted into multiple transformer blocks. Each adapter applies a bottleneck structure that projects the patch features, attends to the prior \mathbf{p} , and reconstructs the output:

$$\begin{aligned} X'_i &= \text{CrossAttn}(\text{Proj}_{\text{down}}(X_i), \mathbf{p}, \mathbf{p}), \\ X_{i+1} &= X_i + \text{Proj}_{\text{up}}(X'_i), \end{aligned} \quad (1)$$

where $X_i \in \mathbb{R}^{N \times d_{\text{up}}}$ is the image feature map at layer i . The final image feature map $\mathbf{X} \in \mathbb{R}^{N \times d}$ is obtained from the last transformer layer with projection. Given the image features and the boxes for the human (b_h), object (b_o), and union region (b_u), we extract region features via RoIAlign [44]:

$$\mathbf{x}_h = \text{RoIAlign}(\mathbf{X}, b_h), \quad \mathbf{x}_o = \text{RoIAlign}(\mathbf{X}, b_o), \quad \mathbf{x}_{\bar{u}} = \text{RoIAlign}(\mathbf{X}, b_u), \quad (2)$$

where $\mathbf{x}_h, \mathbf{x}_o, \mathbf{x}_{\bar{u}} \in \mathbb{R}^d$ denote the region features for the human, object, and union regions, respectively. To enhance the union-region representation with spatial priors, we apply a spatial head that fuses $\mathbf{x}_{\bar{u}}$ with human and object features, as well as their bounding boxes:

$$\mathbf{x}_u = \text{SpatialHead}(\mathbf{x}_{\bar{u}}; \mathbf{x}_h, \mathbf{x}_o; b_h, b_o) \in \mathbb{R}^d. \quad (3)$$

More details on $\text{SpatialHead}(\cdot)$ are provided in the supplementary material. Next, each region feature— \mathbf{x}_h , \mathbf{x}_o , and \mathbf{x}_u —is matched against a set of verb prompts.

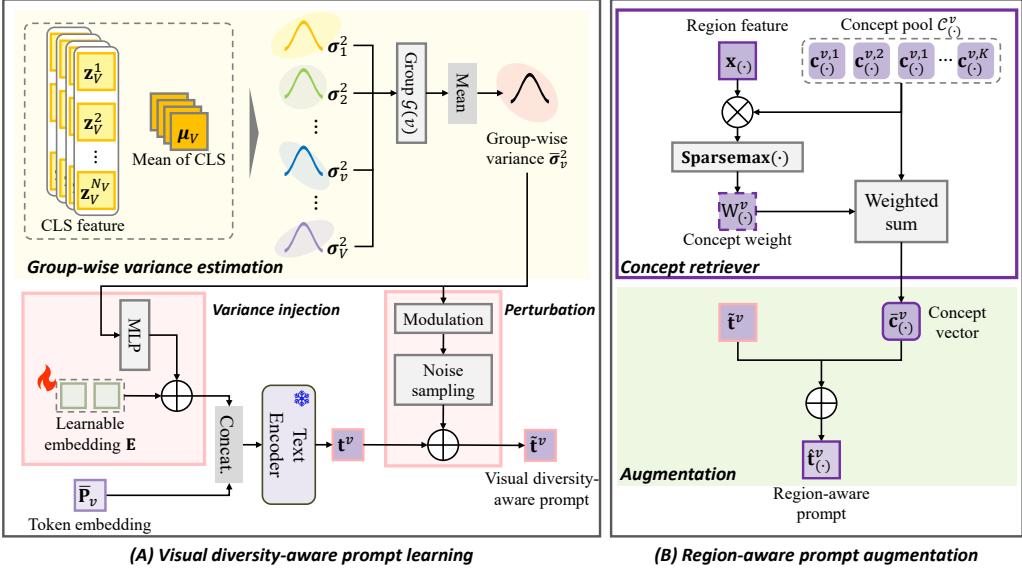


Figure 3: **Detailed architecture of our methods.** (A) To model intra-class variation, we compute verb-wise visual variance σ_v^2 from union-region CLS features, average them over similar verbs to obtain group-wise variance $\bar{\sigma}_v^2$, and inject it into the shared context embedding \mathbf{E} via an MLP. This is combined with the verb prompt $\bar{\mathbf{P}}_v$ and encoded by the CLIP text encoder to produce \mathbf{t}^v , which is further perturbed using Gaussian noise scaled by visual variance. (B) For inter-class discriminability, we retrieve region concepts from features $\mathbf{x}_{(.)}$ using a Sparsemax over a concept pool $\mathcal{C}_{(.)}^v$, and add the result to $\tilde{\mathbf{t}}^v$ to obtain the final region-aware prompt $\hat{\mathbf{t}}_{(.)}^v$.

As illustrated in Fig.2-(B) and (C), Each prompt $\hat{\mathbf{t}}_{(.)}^v$ is constructed in two stages. Let $\mathbf{T}_{(.)} = [\hat{\mathbf{t}}_{(.)}^1, \dots, \hat{\mathbf{t}}_{(.)}^V] \in \mathbb{R}^{d \times V}$ denote the region-aware prompts for $(.) \in \{h, o, u\}$. We first generate a visual diversity-aware prompt $\tilde{\mathbf{t}}^v$ via group-wise variance injection and Gaussian perturbation, then augment it with region-specific concepts retrieved based on region feature $\mathbf{x}_{(.)}$. This two-step design allows the final region-aware prompts to reflect both the visual diversity of verbs and localized context. See Sections 3.2 and 3.3 for details. The logits for each region prompts is computed as:

$$\text{Logit}_h = \mathbf{x}_h^\top \mathbf{T}_h, \quad \text{Logit}_o = \mathbf{x}_o^\top \mathbf{T}_o, \quad \text{Logit}_u = \mathbf{x}_u^\top \mathbf{T}_u, \quad (4)$$

where $\text{Logit}_{(.)} \in \mathbb{R}^V$. Finally, the overall HOI classification logit is obtained by averaging region-wise logits:

$$\text{Logit}_{\text{hoi}} = \frac{1}{3} (\text{Logit}_h + \text{Logit}_o + \text{Logit}_u). \quad (5)$$

We train the model using focal loss [45] for multi-label verb classification.

3.2 Visual diversity-aware prompt learning

To address intra-class visual diversity in HOI detection, we propose a visual diversity-aware prompt learning method that incorporates visual variance into both the context modulation and prompt perturbation processes. Static prompts—optimized as single points—struggle to represent the wide variation of instances within the same verb class. Inspired by recent findings [40, 37, 36] showing that variance-aware representations improve generalization, we explicitly encode class-level visual variance to guide prompt adaptation and inject noise that reflects the extent of such diversity.

Group-wise variance estimation. As illustrated in Fig. 3-(A), we begin by extracting union-region features from the training set using a frozen CLIP image encoder. Specifically, we crop the union box of each human-object pair and extract the CLS token. For each verb v , let $\mathbf{z}_v^{(j)}$ denote the CLS feature of the j -th instance. We compute the mean μ_v and variance σ_v^2 over all N_v samples. To obtain a stable estimate for rare or unseen verbs, we construct a group of verbs $\mathcal{G}(v)$ by selecting the

similar verbs based on cosine similarity between CLIP text embeddings. This grouping allows each verb to inherit variance statistics from its semantically similar neighbors. The group-wise variance is then computed as:

$$\bar{\sigma}_v^2 = \frac{1}{|\mathcal{G}(v)|} \sum_{v' \in \mathcal{G}(v)} \sigma_{v'}^2. \quad (6)$$

This group-wise variance serves as an inductive prior capturing the expected diversity of verb v .

Visual diversity-aware prompts. Then, we transform the group-wise variance into a modulation vector using a lightweight MLP to perform variance injection as follows:

$$\mathbf{d}_v = \text{MLP}(\bar{\sigma}_v^2) \in \mathbb{R}^d. \quad (7)$$

This is added to the shared context embedding $\mathbf{E} \in \mathbb{R}^{N_{\text{ctx}} \times d}$ to produce a verb-specific context:

$$\hat{\mathbf{E}}_v = \mathbf{E} + \mathbf{d}_v \alpha, \quad (8)$$

where α is a small scaling factor for stability. Given a verb prompt sentence \mathbf{P}_v (e.g., “A photo of a person is [v]+ing an object.”), we tokenize it into token embedding $\bar{\mathbf{P}}_v$ and concatenate it with the modulated context to form the final input for CLIP text encoder $\mathcal{T}(\cdot)$:

$$\mathbf{t}^v = \mathcal{T}([\hat{\mathbf{E}}_v; \bar{\mathbf{P}}_v]) \in \mathbb{R}^d. \quad (9)$$

To further reflect visual variability, we perturb each prompt embedding using Gaussian noise scaled by the group-wise standard deviation. We normalize $\bar{\sigma}_v$ across dimensions and modulate it to match the standard deviation of \mathbf{t}^v , producing $\tilde{\sigma}_v$. Noise is then sampled and applied element-wise:

$$\tilde{\mathbf{t}}^v = \mathbf{t}^v + (\epsilon \odot \tilde{\sigma}_v) \beta, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (10)$$

where β is a small scaling factor controlling the perturbation strength. This results in prompt embeddings that encode both the central semantics of the verb and its expected visual diversity. Finally, we collect all perturbed prompts across the V verbs to form the diversity-aware prompts $\tilde{\mathbf{T}} = [\tilde{\mathbf{t}}^1, \dots, \tilde{\mathbf{t}}^V] \in \mathbb{R}^{d \times V}$, which serves as the base for region-aware prompts.

3.3 Region-aware prompt augmentation

To address inter-class visual entanglement—where semantically distinct verbs exhibit similar visual patterns—we augment prompts using region-specific concepts from the human, object, and union regions. While diversity-aware prompts reflect class-level variation, they cannot capture the region concepts needed to distinguish visually similar verbs. To bridge this gap, we introduce retrieval-based region-aware prompt augmentation, which allows each prompt to specialize based on regions.

Region concept generation. To enrich prompt embeddings with localized semantics, we follow [46, 47, 20] and query LLMs (e.g., LLaMA-7B [48] and ChatGPT-4 [49]) to generate K region-level concepts for each verb v and region type $\mathbf{R} \in \{\text{human, object, union}\}$. Each prompt follows the format: “For the verb $[\mathbf{P}_v]$, give K short visual concepts for the $[\mathbf{R}]$ region.” The resulting concepts are encoded using the CLIP text encoder $\mathcal{T}(\cdot)$ to form the concept pool:

$$\mathcal{C}_{(\cdot)}^v = \left\{ \mathbf{c}_{(\cdot)}^{v,1}, \dots, \mathbf{c}_{(\cdot)}^{v,K} \right\}, \quad \mathbf{c}_{(\cdot)}^{v,k} \in \mathbb{R}^d,$$

where $(\cdot) \in \{\text{h, o, u}\}$ denotes the region. For more detailed examples, please refer to Fig.4-(A).

Region-aware prompts. As shown in Fig. 3-(B), given a region feature $\mathbf{x}_{(\cdot)} \in \mathbb{R}^d$ and its corresponding concept pool $\mathcal{C}_{(\cdot)}^v$, we compute cosine similarity between the feature and each concept:

$$s_{(\cdot)}^{v,k} = \frac{\langle \mathbf{x}_{(\cdot)}, \mathbf{c}_{(\cdot)}^{v,k} \rangle}{\|\mathbf{x}_{(\cdot)}\| \|\mathbf{c}_{(\cdot)}^{v,k}\|}. \quad (11)$$

To highlight the most informative concepts while ignoring irrelevant ones, we apply Sparsemax [50] to the similarity scores, which assigns exact zero weights to uninformative entries and retains only the most relevant concepts. Using the resulting scalar weight $W_{(\cdot)}^{v,k}$, we compute a region concept vector:

$$\bar{\mathbf{c}}_{(\cdot)}^v = \sum_{k=1}^K \mathbf{c}_{(\cdot)}^{v,k} W_{(\cdot)}^{v,k} \in \mathbb{R}^d. \quad (12)$$

Table 1: Comparison under NF and RF settings. We report harmonic mean (HM) between Unseen and Seen. Best in **bold**, second best underlined.

Method	Backbone	NF-UC				RF-UC			
		HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
GEN-VLKT [23]	Resnet50+ViT-B	24.17	23.71	25.05	23.38	26.08	30.56	21.36	32.91
EoID [52]	Resnet50	26.71	26.69	26.76	26.66	26.11	29.52	22.04	31.39
HOICLIP [27]	Resnet50+ViT-B	28.70	27.75	29.36	28.10	26.55	32.99	25.83	28.47
ADA-CM [10]	Resnet50+ViT-B	<u>31.76</u>	<u>31.39</u>	32.41	<u>31.13</u>	30.48	33.01	27.63	34.35
CLIP4HOI [26]	Resnet50+ViT-B	29.54	28.90	31.44	28.26	<u>31.23</u>	34.08	27.88	35.48
CMMMP [24]	Resnet50+ViT-B	30.82	30.18	32.09	29.71	31.10	32.18	<u>29.45</u>	32.87
EZ-HOI [25]	Resnet50+ViT-B	31.76	31.17	<u>33.66</u>	30.55	31.18	33.13	29.02	34.15
Ours	Resnet50+ViT-B	33.85	32.57	36.45	31.60	32.77	<u>33.78</u>	31.29	<u>34.41</u>

This concept vector is then used to augment the diversity-aware prompt \mathbf{t}^v , producing the final region-aware prompt as follows:

$$\hat{\mathbf{t}}_{(\cdot)}^v = \mathbf{t}^v + \bar{\mathbf{c}}_{(\cdot)}^v \gamma \in \mathbb{R}^d, \quad (13)$$

where γ is a scalar controlling the degree of augmentation. The final set of region-aware prompts $\mathbf{T}_{(\cdot)} = [\hat{\mathbf{t}}_{(\cdot)}^1, \dots, \hat{\mathbf{t}}_{(\cdot)}^V]$ is used to compute classification logits, as described in Eq. (4). This improves discriminability of the model by aligning each prompt with region-specific semantics.

4 Experiments

4.1 Experimental settings

Datasets. We conduct experiments on the HICO-DET benchmark for HOI detection. HICO-DET contains 80 object categories from the COCO dataset [51] and 117 actions, forming 600 HOI classes. It includes 47,776 images, with 38,118 for training and 9,658 for testing.

Zero-shot setting on HICO-DET. Following prior works [3, 2, 1, 23], we evaluate under four settings: Non-rare First Unseen Composition (NF-UC), Rare First (RF-UC), Unseen Object (UO), and Unseen Verb (UV). NF-UC and RF-UC define 120 unseen and 480 seen HOI triplets from 600 total, with unseen compositions drawn from head and tail categories, respectively. UO uses 68 object classes to construct 500 seen and 100 unseen triplets. UV withholds 20 out of 117 verb classes, yielding 516 seen and 84 unseen triplets.

Evaluation metric. Mean Average Precision (mAP) is used to evaluate the model for HOI detection. Specifically, a sample is regarded as a true positive if two conditions are met: 1) the IoU of both human and object bounding boxes is larger than 0.5, and 2) the HOI triplet prediction is correct.

Implementation Details. We follow the standard training setup used in prior zero-shot two-stage HOI detection methods [24, 10, 25], where DETR is first fine-tuned on instance-level annotations from the HICO-DET training split. Unless otherwise noted, we use CLIP ViT-B/16 as the visual backbone. Additional implementation and training details are provided in the supplementary material.

4.2 Zero-shot HOI detection

We evaluate our method on four zero-shot HOI detection settings in HICO-DET: NF-UC, RF-UC, UO, and UV. Table. 1, 2, and 3 report mAP (Full / Unseen / Seen), harmonic mean (HM), and trainable parameters (#TP), comparing with state-of-the-art baselines. In NF-UC and RF-UC (Table. 1), our method achieves the best scores across all metrics. In NF-UC, we obtain 36.45 (Unseen) and 33.85 (HM), surpassing CLIP4HOI by **+5.01** (Unseen) and **+4.31** (HM). In RF-UC, we outperform EZ-HOI by **+2.27** (Unseen) and **+1.59** (HM). Our method uses only **4.50M** trainable parameters, much fewer than EZ-HOI (6.85M) and CLIP4HOI (56.7M). In UO (Table. 2), we achieve **36.13** (Unseen) and **34.41** (HM), outperforming CMMMP and EZ-HOI by **+1.97** and **+2.27** HM, respectively. In UV (Table. 3), we again achieve the best scores: **26.69** (Unseen), **32.72** (Full), and **29.80** (HM), with a **+1.59** gain on Unseen verbs over EZ-HOI. These results confirm the effectiveness of our *visual diversity-aware prompts* and *region-aware prompts* in the zero-shot HOI detection.

Table 2: Comparison under the UO (Unseen Object) setting. We report harmonic mean (HM) between Unseen and Seen. #TP: trainable parameters. Best in **bold**, second best underlined.

Method	Setting	Backbone	#TP	HM	Full	Unseen	Seen
FCL [1]	UO	Resnet50	–	17.65	19.87	15.54	20.74
ATL [2]	UO	Resnet50	–	17.79	20.47	15.11	21.54
GEN-VLKT [23]	UO	Resnet50	42.05M	20.11	25.63	15.01	28.92
HOICLIP [27]	UO	Resnet50+ViT-B	66.18M	20.32	28.53	16.30	30.99
CLIP4HOI [26]	UO	Resnet50+ViT-B	56.7M	31.98	<u>32.58</u>	31.79	<u>32.73</u>
CMMP [24]	UO	Resnet50+ViT-B	2.30M	<u>32.44</u>	31.59	<u>33.76</u>	31.15
EZ-HOI [25]	UO	Resnet50+ViT-B	6.85M	32.14	32.27	33.28	32.06
Ours	UO	Resnet50+ViT-B	4.50M	34.41	33.39	36.13	32.84

Table 3: Comparison under the UV (Unseen Verb) setting. We report harmonic mean (HM) between Unseen and Seen. #TP: trainable parameters. Best in **bold**, second best underlined.

Method	Setting	Backbone	#TP	HM	Full	Unseen	Seen
GEN-VLKT [23]	UV	Resnet50+ViT-B	42.05M	24.35	28.74	20.96	30.23
EoID [52]	UV	Resnet50	–	26.29	29.61	22.71	30.73
HOICLIP [27]	UV	Resnet50+ViT-B	66.18M	27.72	31.09	24.30	32.19
CLIP4HOI [26]	UV	Resnet50+ViT-B	56.7M	28.35	30.42	26.02	31.14
CMMP [24]	UV	Resnet50+ViT-B	2.30M	<u>29.23</u>	31.84	<u>26.23</u>	32.75
EZ-HOI [25]	UV	Resnet50+ViT-B	6.85M	29.09	<u>32.32</u>	25.10	<u>33.49</u>
Ours	UV	Resnet50+ViT-B	4.50M	29.80	32.73	26.69	33.72

Table 4: **Ablation results under four zero-shot settings.** VDP: visual diversity-aware prompts, RAP: region-aware prompts, VDRP: full model with both.

NF-UC	Full	Unseen	Seen	RF-UC	Full	Unseen	Seen
BASE	28.99	31.68	28.32	BASE	29.80	25.64	30.84
+ VDP	30.17	32.19	29.66	+ VDP	31.95	29.16	32.65
+ RAP	30.43	34.93	29.30	+ RAP	32.43	26.46	33.93
+ VDRP	32.57	36.45	31.60	+ VDRP	33.78	31.29	34.41
UO	Full	Unseen	Seen	UV	Full	Unseen	Seen
BASE	28.92	28.60	30.50	BASE	29.72	22.41	30.91
+ VDP	31.49	33.29	31.13	+ VDP	31.53	23.78	32.79
+ RAP	31.85	33.90	31.44	+ RAP	31.28	24.53	32.38
+ VDRP	33.39	36.13	32.84	+ VDRP	32.73	26.69	33.72

4.3 Ablation studies

Component analysis. To assess each component’s contribution, we conduct an ablation study across four zero-shot HOI settings (Table 4). Starting from a static prompt baseline, we add two modules: (1) **VDP** models intra-class variation using group-wise visual variance and Gaussian noise, and (2) **RAP** augments prompt embeddings with retrieved region-specific concepts. **VDP** improves generalization by capturing the variance of diverse verb appearances, while **RAP** enhances discrimination among visually similar verbs. The full model **VDRP** achieves the best performance across all settings, demonstrating the complementary benefits of both modules.

Group-wise modeling in VDP. To evaluate the impact of group-wise visual variance, we compare different configurations in the **VDP** module based on the number of verbs per group: (1) a single global variance shared across all verbs (Global), (2) group-wise variance using clusters of 3 verbs each (Group_3), and (3) clusters of 5 verbs each (Group_5). As shown in Table 5, we observe that group-wise variance modeling generally improves over a global estimate, but performance varies depending on the group size.

Impact of injection scale. We evaluate variance injection scale $\alpha \in \{0.01, 0.02, 0.10\}$ to analyze its influence on the delta scaling in Eq. 8. Results across all settings are summarized in Table 6. We

Table 5: Impact of group configuration in visual diversity-aware prompts (VDP). Best results in **bold**.

Method	NF-UC			RF-UC			UO			UV		
	Full	Unseen	Seen									
Global	31.92	35.92	30.92	33.94	32.27	34.36	32.30	34.23	31.91	32.55	26.07	33.60
Group_3	32.43	35.89	31.57	33.15	29.57	34.04	32.70	34.76	32.29	32.75	26.25	33.81
Group_5	32.57	36.45	31.60	33.78	31.29	34.41	33.39	36.13	32.84	32.73	26.69	33.72

Table 6: Impact of injection scale α on performance across zero-shot settings. Best results in **bold**.

α	NF-UC			RF-UC			UO			UV		
	Full	Unseen	Seen									
0.01	32.03	35.29	31.22	33.51	30.58	34.24	32.81	35.52	32.27	32.48	25.97	33.54
0.10	32.02	35.80	31.07	33.16	29.86	33.98	32.81	35.11	32.35	32.54	24.73	33.81
0.02	32.57	36.45	31.60	33.78	31.29	34.41	33.39	36.13	32.84	32.73	26.69	33.72

Table 7: Effect of Gaussian perturbation across four zero-shot settings. Best results are in **bold**.

Method	NF-UC				RF-UC			
	HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
W/O perturbation	33.57	32.49	35.68	31.69	31.76	33.12	29.83	33.95
W/ perturbation	33.85	32.57	36.45	31.60	32.77	33.78	31.29	34.41
Method								
Method	UO				UV			
	HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
W/O perturbation	33.67	33.09	34.59	32.79	29.49	32.72	26.16	33.79
W/ perturbation	34.41	33.39	36.13	32.84	29.80	32.73	26.69	33.72

observe that $\alpha = 0.02$ consistently achieves the best trade-off between unseen and seen performance. This value was chosen to match the initialization scale of CLIP’s context embeddings, ensuring stability as excessively large α causes over-perturbed context tokens and unstable training. Given that context tokens are highly sensitive to initialization, keeping the scale aligned promotes stable optimization while maintaining sufficient diversity.

Effect of Gaussian perturbation. We study the effect of Gaussian perturbation on prompt embeddings by varying the scale $\beta \in \{0, 0.1\}$. Table 7 summarizes the results across all zero-shot settings. Perturbation improves generalization in most unseen cases, confirming that stochastic sampling helps the model capture diverse visual realizations of each verb.

Retrieval strategy. We compare three retrieval methods for region concepts: Softmax, Top-3, and Sparsemax (Table 8). Overall, Sparsemax performs most consistently, benefiting from its ability to suppress irrelevant concepts and emphasize informative ones.

Impact of augmentation scaling factor. We evaluate $\gamma \in \{0.2, 0.5, 1.0\}$ in Eq. 13 to study how strongly region-derived cues contribute to prompt adaptation. Results are reported in Table 9. A moderate scaling of $\gamma = 0.2$ provides the best overall balance, whereas excessive cue weighting ($\gamma = 1.0$) slightly degrades performance across multiple settings, likely due to noisy region concepts.

4.4 Qualitative results

Fig. 4 shows qualitative examples of region-specific concept generation and retrieval. Given a verb prompt (e.g., “licking an object”), our method uses LLMs (e.g., LLaMA-7B [48] and GPT-4 [49]) to generate K region concepts, capturing fine-grained concepts such as facial expression or visual patterns. While LLaMA-7B provides diverse concept descriptions, we observed noise and redundancy in object-related concepts; therefore, GPT-4 was additionally used (see Supplementary for details). As illustrated, similar verbs like “licking” and “eating” are associated with distinct region-specific signals—enabling the model to better distinguish subtle visual differences, enhancing the discriminability.

Table 8: Comparison of retrieval strategies for region-aware prompts (RAP). Best results in **bold**.

Method	NF-UC			RF-UC			UO			UV		
	Full	Unseen	Seen									
Softmax(\cdot)	32.01	35.93	31.03	33.32	30.37	34.06	32.91	35.83	32.32	32.57	25.45	33.74
Top-3	31.97	35.85	31.00	33.35	30.80	33.99	32.94	33.57	32.41	32.47	25.85	33.55
Sparsemax(\cdot)	32.57	36.45	31.60	33.78	31.29	34.41	33.39	36.13	32.84	32.73	26.69	33.72

Table 9: Effect of augmentation scale γ on performance across all settings. Best results in **bold**.

γ	NF-UC			RF-UC			UO			UV		
	Full	Unseen	Seen									
0.5	32.14	35.92	31.19	33.52	30.73	34.21	33.08	35.17	32.66	32.72	27.06	33.64
1.0	31.85	35.15	31.02	33.48	30.67	34.18	32.60	34.74	32.17	32.54	25.92	33.62
0.2	32.57	36.45	31.60	33.78	31.29	34.41	33.39	36.13	32.84	32.73	26.69	33.72

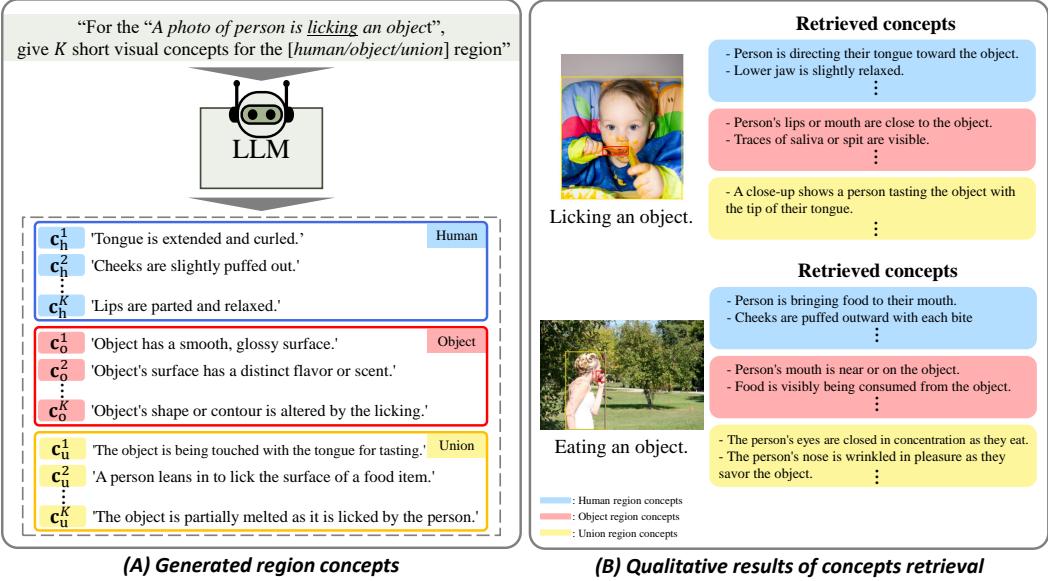


Figure 4: **Qualitative examples of region concept generation and retrieval.** (A) Given a verb prompt and region type, an LLM generates K region concepts per verb. (B) Retrieved concepts for “Licking” and “Eating” highlight subtle region concepts that help disambiguate visually similar interactions. Concepts are color-coded by region: blue (human), red (object), yellow (union).

5 Conclusion

We present **VDRP**, a prompt learning framework for zero-shot HOI detection that addresses the *visual complexity of interaction*, including intra-class diversity and inter-class entanglement among verb classes. To this end, we propose a *visual diversity-aware prompt learning* that injects group-wise visual variance into the prompt context and applies Gaussian perturbation, enabling prompt embeddings to better reflect diverse visual appearances. We also introduce *region-aware prompt augmentation*, which augments diversity-aware prompts with region-specific concepts from the human, object, and union regions to improve verb discriminability. Extensive experiments on HICO-DET show that our method outperforms prior work across four zero-shot settings while maintaining high parameter efficiency. These results underscore the value of combining distributional modeling and region-level augmentation for generalizable zero-shot HOI detection.

6 Acknowledgments

This work was partly supported by Korea Research Institute for defense Technology planning and advancement - Grant funded by Defense Acquisition Program Administration (DAPA) (KRIT-CT-23-021, 45%), by the Institute of Information & Communications Technology Planning & Evaluation (IITP)-ITRC (Information Technology Research Center) grant funded by the Korea government (MSIT) (IITP-2025-RS-2024-00436857, 45%) and by the National Research Foundation of Korea (RS-2022-NR068758, 10%). We thank **Jaewoon Byun** for helpful discussion.

References

- [1] Zhi Hou, Baosheng Yu, Yu Qiao, Xiaojiang Peng, and Dacheng Tao. Detecting human-object interaction via fabricated compositional learning. In *CVPR*, 2021.
- [2] Zhi Hou, Baosheng Yu, Yu Qiao, Xiaojiang Peng, and Dacheng Tao. Affordance transfer learning for human-object interaction detection. In *CVPR*, 2021.
- [3] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. In *ECCV*, 2020.
- [4] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jiashi Feng. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *CVPR*, 2020.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] Ting Lei, Fabian Caba, Qingchao Chen, Hailin Jin, Yuxin Peng, and Yang Liu. Efficient adaptive human-object interaction detection with concept-guided memory. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6480–6490, 2023.
- [11] Yichao Cao, Qingfei Tang, Xiu Su, Song Chen, Shan You, Xiaobo Lu, and Chang Xu. Detecting any human-object interaction relationship: Universal hoi detector with spatial prompt learning on foundation models. *Advances in Neural Information Processing Systems*, 36:739–751, 2023.
- [12] Guangzhi Wang, Yangyang Guo, Ziwei Xu, and Mohan Kankanhalli. Bilateral adaptation for human-object interaction detection with occlusion-robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27970–27980, 2024.
- [13] Jihwan Park, SeungJun Lee, Hwan Heo, Hyeong Kyu Choi, and Hyunwoo J Kim. Consistency learning via decoding path augmentation for transformers in human object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1019–1028, 2022.
- [14] 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 *ICML*, 2021.

[15] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022.

[16] Sehyung Kim, Chanhyeong Yang, Jihwan Park, Taehoon Song, and Hyunwoo J Kim. Super-class guided transformer for zero-shot attribute classification. *arXiv preprint arXiv:2501.05728*, 2025.

[17] Mingfei Gao, Chen Xing, Juan Carlos Niebles, Junnan Li, Ran Xu, Wenhao Liu, and Caiming Xiong. Open vocabulary object detection with pseudo bounding-box labels. In *ECCV*, 2022.

[18] 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.

[19] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16816–16825, 2022.

[20] Jooyeon Kim, Eulrang Cho, Sehyung Kim, and Hyunwoo J Kim. Retrieval-augmented open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17427–17436, 2024.

[21] Huaizu Jiang, Xiaojian Ma, Weili Nie, Zhiding Yu, Yuke Zhu, and Anima Anandkumar. Bongard-hoi: Benchmarking few-shot visual reasoning for human-object interactions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19056–19065, 2022.

[22] Jie Yang, Bingliang Li, Ailing Zeng, Lei Zhang, and Ruimao Zhang. Open-world human-object interaction detection via multi-modal prompts. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition*, pages 16954–16964, 2024.

[23] 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 *CVPR*, 2022.

[24] Ting Lei, Shaofeng Yin, Yuxin Peng, and Yang Liu. Exploring conditional multi-modal prompts for zero-shot hoi detection. In *European Conference on Computer Vision*, pages 1–19. Springer, 2024.

[25] Qinjian Lei, Bo Wang, and Robby Tan. Ez-hoi: Vlm adaptation via guided prompt learning for zero-shot hoi detection. *Advances in Neural Information Processing Systems*, 37:55831–55857, 2024.

[26] Yunyao Mao, Jiajun Deng, Wengang Zhou, Li Li, Yao Fang, and Houqiang Li. Clip4hoi: Towards adapting clip for practical zero-shot hoi detection. *Advances in Neural Information Processing Systems*, 36:45895–45906, 2023.

[27] Shan Ning, Longtian Qiu, Yongfei Liu, and Xuming He. Hoiclip: Efficient knowledge transfer for hoi detection with vision-language models. In *CVPR*, 2023.

[28] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, 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*, 2020.

[29] Bingjie Xu, Yongkang Wong, Junnan Li, Qi Zhao, and Mohan S Kankanhalli. Learning to detect human-object interactions with knowledge. In *CVPR*, 2019.

[30] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.

[31] Jongha Kim, Jihwan Park, Jinyoung Park, Jinyoung Kim, Sehyung Kim, and Hyunwoo J Kim. Groupwise query specialization and quality-aware multi-assignment for transformer-based visual relationship detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 28160–28169, 2024.

[32] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023.

[33] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022.

[34] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021.

[35] Jinguo Luo, Weihong Ren, Weibo Jiang, Xi’ai Chen, Qiang Wang, Zhi Han, and Honghai Liu. Discovering syntactic interaction clues for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28212–28222, 2024.

[36] Eulrang Cho, Jooyeon Kim, and Hyunwoo J Kim. Distribution-aware prompt tuning for vision-language models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22004–22013, 2023.

[37] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5206–5215, 2022.

[38] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fa-had Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19113–19122, 2023.

[39] Wentao Bao, Lichang Chen, Heng Huang, and Yu Kong. Prompting language-informed distribution for compositional zero-shot learning. In *European Conference on Computer Vision*, pages 107–123. Springer, 2024.

[40] Xingyu Zhu, Beier Zhu, Yi Tan, Shuo Wang, Yanbin Hao, and Hanwang Zhang. Enhancing zero-shot vision models by label-free prompt distribution learning and bias correcting. *Advances in Neural Information Processing Systems*, 37:2001–2025, 2024.

[41] Jinyoung Park, Juyeon Ko, and Hyunwoo J Kim. Prompt learning via meta-regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26940–26950, 2024.

[42] Dongjun Lee, Seokwon Song, Jihee Suh, Joonmyeong Choi, Sanghyeok Lee, and Hyunwoo J Kim. Read-only prompt optimization for vision-language few-shot learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1401–1411, 2023.

[43] Songlin Dong, Zhengdong Zhou, Chenhao Ding, Xinyuan Gao, Alex Kot, and Yihong Gong. Diversity covariance-aware prompt learning for vision-language models. *arXiv preprint arXiv:2503.01531*, 2025.

[44] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

[45] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.

[46] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *International conference on machine learning*, pages 5338–5348. PMLR, 2020.

[47] Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19187–19197, 2023.

[48] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

[49] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[50] Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International conference on machine learning*, pages 1614–1623. PMLR, 2016.

[51] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.

[52] Mingrui Wu, Jiaxin Gu, Yunhang Shen, Mingbao Lin, Chao Chen, and Xiaoshuai Sun. End-to-end zero-shot hoi detection via vision and language knowledge distillation. In *Proceedings of the AAAI Conference on artificial intelligence*, volume 37, pages 2839–2846, 2023.

[53] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[54] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[55] 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, October 2021.

[56] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015.

[57] 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.

[58] 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*, pages 21614–21624, 2023.

[59] Liulei Li, Jianan Wei, Wenguan Wang, and Yi Yang. Neural-logic human-object interaction detection. *Advances in Neural Information Processing Systems*, 36:21158–21171, 2023.

[60] Sedigheh Eslami and Gerard de Melo. Mitigate the gap: Improving cross-modal alignment in CLIP. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=aPTGvFqile>.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We clearly state the main contributions and scope in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss the limitations of our method—such as the Gaussian assumption in visual space and the reliance on LLM-generated concepts—in the supplemental material.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our work does not include any theoretical results or formal proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide implementation details and training configurations in the supplementary material to ensure reproducibility of all main experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We will release our code and data upon publication, along with instructions in the supplementary material to reproduce all main results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We describe the training and evaluation protocols in Section 4 and provide additional details on implementation and hyperparameters in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[NA\]](#)

Justification: Our evaluation follows standard zero-shot HOI detection benchmarks with deterministic protocols, so statistical significance is not reported.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We report training time and GPU specifications in the supplementary material to support reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We confirm that our research adheres to the NeurIPS Code of Ethics, including responsible data usage and transparency in methodology and reporting.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: We briefly discuss societal benefits and risks in the conclusion under “Societal Impact.”

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work does not involve models or data with high risk of misuse, so this question is not applicable.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We use publicly available datasets (HICO-DET) and pretrained models (e.g., CLIP, DETR) that are properly cited in the paper, and their licenses and usage terms are fully respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We properly credit and cite all existing datasets, models (e.g., CLIP, HICO-DET), and related assets used in our work, and follow their license terms as described in the paper and supplementary material.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our work does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our study does not involve human subjects and therefore does not require IRB approval.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [\[Yes\]](#)

Justification: We use a large language model to generate region-level concepts (e.g., LLaMA-7B) as part of our retrieval-based prompt augmentation module.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Appendix

A Implementation details

We conduct all experiments using PyTorch with mixed-precision training. For the base model using CLIP ViT-B/16, we train on two NVIDIA GeForce RTX 3090 GPUs with a batch size of 8 for 12 epochs. We use the AdamW [53] optimizer with an initial learning rate of 1×10^{-3} , decayed to 1×10^{-4} using a cosine scheduler and a weight decay of 8. We report the best results after applying weight decay. The framework adopts CLIP ViT-B/16 as the visual encoder, where the adapter applies a bottleneck transformation from $d_{\text{up}} = 768$ to $d_{\text{down}} = 64$, followed by projection into a shared embedding space of dimension $d = 512$. Human-object detection is performed by a frozen DETR model with Resnet50 [54], and detections with confidence below the threshold $\theta = 0.2$ are discarded. The context embedding consists of $N_{\text{ctx}} = 24$ learnable tokens, each initialized from a Gaussian distribution with standard deviation 0.02. To incorporate group-wise variance, we inject a modulation vector (scaled by $\alpha = 0.02$) into the context embedding and apply Gaussian perturbation to the resulting prompt embedding with a noise scale $\beta = 0.1$. For region-aware prompt augmentation, we generate $K = 10$ region concepts per verb and region type (human, object, union) using LLaMA-7B [48] and ChatGPT-4 [49]. These are encoded by the frozen CLIP text encoder and aggregated via Sparsemax [50] to form a concept vector, which is added to the prompt embedding with scaling factor $\gamma = 0.2$. We additionally conduct experiments with a scaled-up version using CLIP ViT-L/14 as the visual encoder. In this setting, the adapter transforms $d_{\text{up}} = 1024$ to $d_{\text{down}} = 64$, with the final embedding dimension set to $d = 768$. Due to memory constraints, these experiments are performed on two NVIDIA RTX 6000 Ada Generation GPUs with a reduced batch size of 4 per GPU.

B Details of our architectures

Details of spatial head. To enhance the union-region representation with geometric priors, we design a spatial head that encodes the spatial relationship between the human and object bounding boxes and fuses it with the region features, inspired by prior works [55, 5, 6]. Given a human box $b_h = [x_1, y_1, x_2, y_2]$ and an object box $b_o = [x'_1, y'_1, x'_2, y'_2]$, we compute the center coordinates, widths, and heights:

$$\mathbf{C}_h = \left(\frac{x_1 + x_2}{2}, \frac{y_1 + y_2}{2} \right), \quad \mathbf{C}_o = \left(\frac{x'_1 + x'_2}{2}, \frac{y'_1 + y'_2}{2} \right),$$

$$w_h = x_2 - x_1, \quad h_h = y_2 - y_1, \quad w_o = x'_2 - x'_1, \quad h_o = y'_2 - y'_1.$$

We then extract spatial features including normalized positions, relative box areas, aspect ratios, intersection-over-union (IoU), and direction-aware relative distances:

$$d_x = \frac{|\mathbf{C}_{ox} - \mathbf{C}_{hx}|}{w_h + \epsilon}, \quad d_y = \frac{|\mathbf{C}_{oy} - \mathbf{C}_{hy}|}{h_h + \epsilon}.$$

These values are concatenated and passed through a small feedforward network to produce a spatial encoding vector \mathbf{E}_S . To incorporate this geometric context, we apply a multi-modal fusion module $\phi_{\text{MMF}}(\cdot)$ that combines the human and object features $[\mathbf{x}_h; \mathbf{X}_o]$ with the spatial encoding:

$$\mathbf{f}_S = \phi_{\text{MMF}}([\mathbf{x}_h; \mathbf{x}_o], \mathbf{E}_S).$$

Finally, the spatial query \mathbf{f}_S is concatenated with the union feature $\mathbf{x}_{\bar{u}}$, and passed through a projection layer to obtain the final spatially-enhanced representation:

$$\mathbf{x}_u = \text{Proj}([\mathbf{x}_{\bar{u}}; \mathbf{f}_S]),$$

which is used for verb classification with visual diversity and region-aware prompts.

Details of sparsemax. To select region-level concepts relevant to each human-object interaction, we adopt the Sparsemax [50] activation function in our region-aware prompt augmentation module. Unlike Softmax, which produces dense probability distributions where all entries are non-zero, Sparsemax enables *sparse* selection by assigning exact zeros to irrelevant entries. This property is particularly useful in our setting, as only a subset of the K generated region concepts are semantically aligned with the visual feature.

Formally, given a score vector $\mathbf{s} \in \mathbb{R}^K$ —representing the cosine similarity between a region feature and K region concept embeddings—we define Sparsemax as a projection onto the probability simplex:

$$\text{Sparsemax}(\mathbf{s}) := \arg \min_{\mathbf{a} \in \Delta^K} \|\mathbf{a} - \mathbf{s}\|^2, \quad \Delta^K = \left\{ \mathbf{a} \in \mathbb{R}^K \mid \sum_{k=1}^K a_k = 1, a_k \geq 0 \right\}. \quad (14)$$

This formulation yields a closed-form solution that projects the input vector onto the probability simplex Δ^K , resulting in a sparse vector \mathbf{a} where low-scoring entries are assigned exact zeros. We apply Sparsemax to the similarity scores between each region feature and the corresponding concept pool, allowing the model to focus on the most informative region concepts while ignoring noisy or irrelevant ones.

C Limitations

While our method demonstrates strong performance across multiple zero-shot HOI detection settings, several broader limitations remain. First, the *region-aware prompts* (RAP) module builds upon region-level concepts generated by large language models (LLMs). Although effective in capturing contextual semantics, LLM-based concepts may be noisy or misaligned with visual concepts due to inherent limitations in language–vision grounding. Improving robustness through confidence-aware filtering or vision-aligned concept refinement is a promising direction. Second, our framework builds on prompt learning, which presents structural challenges when applied to a large number of classes or diverse interaction types. Prompt representations can be sensitive to scaling strategies, initialization, and optimization dynamics, making them less stable in settings with limited data or rare class compositions. This reflects a broader limitation of prompt-based models, where generalization can be affected by prompt granularity, representation collapse, or lack of compositional structure. Future directions may include more robust prompt initialization, adaptive scaling mechanisms, or compositional prompt construction to improve stability and generalization.

D Broader impacts

Our work builds upon vision-language models (VLMs) such as CLIP, which are trained on large-scale web-scraped image-text pairs. As a result, our model may inadvertently inherit biases present in the pretraining data, such as cultural stereotypes or imbalanced representations across demographic groups. While our method aims to improve generalization in zero-shot HOI detection, deployment in real-world applications should be done with caution—particularly in contexts involving surveillance, behavioral analysis, or human activity interpretation. Furthermore, the retrieval-based region prompt augmentation using large language models (e.g., LLaMA-7B and ChatGPT-4), which may also reflect biases or hallucinated associations. Misinterpretation of human-object interactions in sensitive domains (e.g., law enforcement or healthcare) could lead to harmful outcomes if such biases are not properly addressed. To mitigate potential misuse, we recommend restricting deployment to controlled environments with human oversight, and further advocate for evaluating fairness metrics across subpopulations when applying the model to downstream applications.

E Further experiments

V-COCO datasets. V-COCO [56] dataset has 80 object categories derived from COCO datasets [51] and includes 29 interactions, resulting 263 HOI compositions. The number of image samples is 10,396 (5,400/4,964 for train/test).

Fully supervised settings. We further evaluate our method under the fully supervised HOI detection setting on HICO-DET and V-COCO, as summarized in Table 10. Among two-stage methods, our model achieves the best performance on HICO-DET, with an mAP of 39.07, surpassing the prior state-of-the-art EZ-HOI [25]. While one-stage methods such as UniHOI and AGER benefit from end-to-end optimization and full supervision of all components, they typically rely on larger networks and longer training schedules. In contrast, our two-stage approach uses fewer learnable parameters and is designed primarily for zero-shot generalization, which likely limits the effectiveness of group-wise variance modeling and hinders the learning of cross-class structure. On

Table 10: State-of-the-art comparison on HICO-DET and V-COCO under the fully-supervised setting. Bold indicates the best-performing method within each group (one-stage vs. two-stage).

Method	HICO-DET			V-COCO	
	Full	Rare	Nonrare	$AP_{\text{role}}^{s_1}$	$AP_{\text{role}}^{s_2}$
<i>One-stage Methods</i>					
GEN-VLKT [23]	33.75	29.25	35.10	62.4	64.5
HOICLIP [27]	34.69	31.12	35.74	63.5	65.0
RLIPV2 [57]	35.38	29.61	37.11	65.9	68.0
AGER [58]	36.75	33.53	37.43	65.7	67.9
LogicHOI [59]	35.47	32.03	36.64	64.6	65.6
UniHOI [11]	40.06	39.91	40.11	65.6	68.3
<i>Two-stage Methods</i>					
UPT [5]	32.62	28.62	33.81	59.0	64.5
ADA-CM [10]	38.40	37.52	38.66	58.6	64.0
CLIP4HOI [26]	35.33	33.95	35.75	—	66.3
CMMMP [24]	38.14	37.75	38.25	—	64.0
EZ-HOI [25]	38.61	37.70	38.89	60.5	66.2
Ours	39.07	39.08	39.06	60.6	66.2

Table 11: Zero-shot HOI detection results under four splits—NF-UC, RF-UC, UO, and UV—using the scaled CLIP (ViT-L). HM is harmonic mean (HM) between Unseen and Seen. Best results are shown in **bold**, and second best are underlined.

Method	NF-UC				RF-UC			
	HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
UniHOI [11]	30.40	31.79	28.45	32.63	30.76	32.27	28.68	33.16
CMMMP [24]	34.50	<u>35.13</u>	33.52	<u>35.53</u>	<u>36.69</u>	<u>37.13</u>	<u>35.98</u>	<u>37.42</u>
EZ-HOI [25]	<u>35.38</u>	34.84	<u>36.33</u>	34.47	35.73	36.73	34.24	37.35
Ours	36.83	36.46	37.48	36.21	37.58	38.13	36.72	38.48

Method	UO				UV			
	HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
UniHOI [11]	25.17	31.56	19.72	34.76	30.50	34.68	26.05	36.78
CMMMP [24]	<u>37.83</u>	<u>36.74</u>	39.67	<u>36.15</u>	<u>33.75</u>	36.38	<u>30.84</u>	37.28
EZ-HOI [25]	37.06	36.38	38.17	36.02	32.84	<u>36.84</u>	28.82	<u>38.15</u>
Ours	38.41	37.81	<u>39.36</u>	37.50	34.31	37.18	31.16	38.16

the V-COCO benchmark, our model achieves 60.6 and 66.2 AP under Scenario 1 and Scenario 2, respectively—comparable to other two-stage models. We attribute this to the small number of verb classes (24 vs. 117 in HICO-DET) and limited dataset size, which reduce the effectiveness of group-wise variance modeling and cross-class structure learning. Overall, these results demonstrate that although our method is designed to address zero-shot HOI detection, it generalizes well to fully supervised settings, retaining strong performance even when supervision is abundant.

Scaled-up CLIP (ViT-L) setting. To assess the effect of scaling the vision backbone, we evaluate our method using the CLIP ViT-L/14 encoder in place of the default ViT-B/16. As shown in Table 11, our model consistently improves performance across all four zero-shot evaluation splits—NF-UC, RF-UC, UO, and UV. Notably, our method achieves the highest harmonic mean (HM) across all settings, reflecting balanced generalization to both seen and unseen verb compositions. While prior methods such as CMMMP [24] and EZ-HOI [25] exhibit competitive results on specific splits, their performance fluctuates across evaluation scenarios. In contrast, our scaled-up model maintains stable gains throughout, demonstrating that visual diversity-aware prompt learning and region-aware prompt augmentation remain effective even when paired with high-capacity vision-language encoders. These results highlight the scalability of our framework even when scaled to stronger backbones.

Table 12: Comparison of using both mean and variance (μ, σ^2) vs. variance-only (σ^2) in the VDP across four zero-shot settings. Best in **bold**.

Setting	Stats	Full	Unseen	Seen
NF-UC	(μ, σ^2)	32.03	35.75	31.10
	σ^2 only	32.57	36.45	31.60
RF-UC	(μ, σ^2)	33.19	30.34	33.91
	σ^2 only	33.78	31.29	34.41
UO	(μ, σ^2)	33.03	34.87	32.66
	σ^2 only	33.39	36.13	32.84
UV	(μ, σ^2)	32.59	25.33	33.77
	σ^2 only	32.73	26.69	33.72

Table 13: Ablation study comparing region branches (Human, Object and Union) and the full model (H+O+U) under different zero-shot settings. Best in **bold**.

Branch	NF-UC				RF-UC			
	HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
Human	33.27	32.03	35.77	31.09	32.19	33.34	30.53	34.04
Object	33.48	32.17	36.16	31.17	31.90	33.32	29.92	34.17
Union	33.26	32.17	35.40	31.37	32.01	33.23	30.28	33.96
H+O+U	33.85	32.57	36.45	31.60	32.77	33.78	31.29	34.41
Branch	UO				UV			
	HM	Full	Unseen	Seen	HM	Full	Unseen	Seen
Human	34.08	32.90	36.11	32.26	29.35	32.66	25.97	33.75
Object	33.39	32.65	34.61	32.26	29.64	33.07	26.15	34.20
Union	33.39	32.70	34.52	32.34	29.55	32.60	26.36	33.62
H+O+U	34.41	33.39	36.13	32.84	29.80	32.73	26.69	33.72

F More ablation studies

Effect of mean vs. variance in VDP. To evaluate the contribution of distributional statistics in our visual diversity-aware prompts (VDP) module, we compare two variants: one that uses only group-wise variance (σ^2), and another that combines both the group-wise mean (μ) and variance (σ^2) through concatenation followed by projection. Table 12 reports results across all four zero-shot settings. Overall, we find that the variance-only variant consistently matches or outperforms the mean-variance combination—especially on the Unseen splits in all settings. For instance, in NF-UC, using variance alone yields the highest unseen mAP of **36.45**, compared to 35.75 with $\mu+\sigma^2$. This suggests that variance alone serves as a stronger signal for modeling intra-class visual diversity, particularly under zero-shot generalization conditions. Prior studies such as Zhu et al. [40] also report that variance-centered prompt representations improve generalizability by capturing the dispersion of visual features without overfitting to sample means. Our results align with this observation, indicating that the inclusion of the mean may introduce redundant or unstable signals—especially for low-shot or noisy verb clusters, where the class prototype is poorly defined. An exception arises in the UV setting, where the combined variant slightly improves performance on seen metrics. Nonetheless, the overall trend supports the effectiveness and robustness of variance-only modeling for capturing distributional diversity in the prompt space.

Effect of branches in RAP. We compare the individual contributions of the human, object, and union region branches under four zero-shot HOI detection settings. The results are summarized in Table 13. Across all settings, the full model consistently outperforms each individual branch, confirming the complementary nature of region-level concepts. Notably, the proposed combination achieves the best unseen mAP in all settings, demonstrating the effectiveness of region-aware prompts in enhancing discriminability. These findings highlight that human and object concepts capture distinct but complementary contextual signals, while the union branch provides holistic spatial grounding that reinforces interaction understanding.

Table 14: Zero-shot verb variance robustness under the UV setting. We remove the top-3 semantic neighbors for each unseen verb before computing visual variance.

	HM	Full	Unseen	Seen
W/O Top-3	29.48	32.64	26.21	33.69
W/ Top-3	29.80	32.73	26.69	33.72

Table 15: Sensitivity analysis of τ -Sparsemax under four zero-shot settings. Best results are in **bold**.

τ value	NF-UC			RF-UC		
	Full	Unseen	Seen	Full	Unseen	Seen
0.10	32.23	35.85	31.33	33.59	30.92	34.26
0.05	32.38	35.76	31.53	33.10	29.56	33.99
0.00	32.57	36.45	31.60	33.78	31.29	34.41

τ value	UO			UV		
	Full	Unseen	Seen	Full	Unseen	Seen
0.10	32.90	35.84	32.31	32.85	26.95	33.81
0.05	32.95	34.63	32.62	32.73	25.91	33.84
0.00	33.39	36.13	32.84	32.73	26.69	33.72

Table 16: Robustness of VDP under few-shot variance sampling ($N_v = 5$). Best results are in **bold**.

NF-UC	Full	Unseen	Seen	RF-UC	Full	Unseen	Seen
Few-shot	32.23	35.90	31.31	Few-shot	33.40	30.59	34.10
All	32.57	36.45	31.60	All	33.78	31.29	34.41

UO	Full	Unseen	Seen	UV	Full	Unseen	Seen
Few-shot	32.55	34.75	32.11	Few-shot	32.89	25.88	34.04
All	33.39	36.13	32.84	All	32.73	26.69	33.72

G Robustness of visual diversity-aware prompt learning

Zero-shot verb variance robustness. To evaluate the robustness of group-wise variance estimation for unseen verbs, we conduct an additional experiment under the UV setting. Specifically, we remove the top-3 semantic neighbors (based on CLIP similarity) for each unseen verb from the training set before computing the group-wise visual variance. This setup simulates a scenario where certain verbs lack semantically related support, allowing us to analyze the stability of variance estimation under reduced contextual guidance. As shown in Tab 14, the removal of neighboring classes results in only a marginal performance drop in unseen verbs (-0.48 mAP). This indicates that the proposed grouping mechanism maintains its generalization ability even when semantic support is limited.

τ -Sparsemax sensitivity. To investigate the effect of sparsity calibration in concept retrieval, we introduce τ -Sparsemax(\cdot), where all values below a threshold $\tau \in \{0.0, 0.05, 0.1\}$ are zeroed post-Sparsemax(\cdot). This modification controls the degree of sparsity in the region-concept selection process. The results in Tab 15 show that excessive pruning ($\tau = 0.05$) slightly decreases overall performance, while higher sparsity ($\tau = 0.10$) recovers seen-class precision at the cost of unseen generalization. The default $\tau = 0.0$ provides the most balanced outcome, indicating that the original Sparsemax(\cdot) formulation already achieves an effective level of adaptive sparsity.

Robustness of visual variance under few-shot sampling. We further test the stability of visual variance estimation when fewer visual samples are available. For this experiment, group-wise covariance matrices are computed using only five randomly sampled features per verb ($N_v = 5$), compared to using all available samples. As summarized in Tab 16, the performance reduction across all settings is minimal (typically less than 0.6 mAP), demonstrating that VDP remains robust even under limited sample diversity.

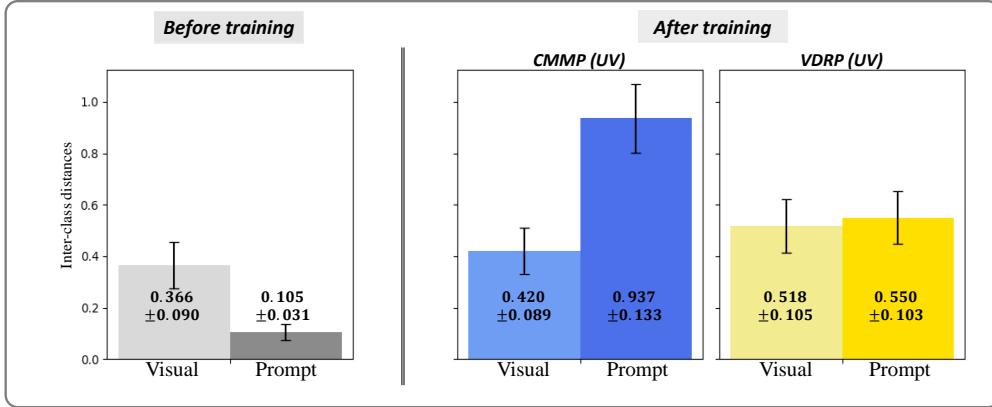


Figure 5: **Pairwise inter-class distances between prompts and visual features.** We report the average pairwise cosine distance (i.e., $D = \mathbb{E}_{i \neq j} [1 - \cos(z_i, z_j)]$) across verb classes, for both visual and prompt embeddings, before and after training. Visual features are extracted from union regions. Before training, we use the CLIP visual encoder applied to cropped union region images. After training, we follow the ROI-Align feature extraction pipeline consistent with the two-stage method (i.e., pooling patch embeddings within the union box). For each verb class, a medoid is selected among all union features to represent its prototype. While prompt embeddings are initially collapsed with low diversity, VDRP maintains a balanced and aligned distribution relative to visual features, unlike CMMP which over-separates prompts and disrupts cross-modal structure.

H Inter-class alignment of visual and prompt representations

To assess inter-class alignment between visual and prompt representations, we analyze the distributional structure of each modality after training. Specifically, we examine whether our visual diversity-aware prompt learning preserves class-level relationships that are consistent with those observed in the visual embedding space. We measure the pairwise cosine distance between class-level prototypes in both visual and prompt embedding spaces:

$$D = \mathbb{E}_{i \neq j} [1 - \cos(z_i, z_j)],$$

where z_i denotes the prototype embedding of the i -th verb class. For the visual side, we extract features from the union region of each verb instance using the CLIP visual encoder. Before training, we use the CLIP visual encoder applied directly to the union-cropped image. After training, we adopt a region-specific pooling approach: patch embeddings within the union ROI-Align [44] are aggregated to match the training architecture of our method. To mitigate over-smoothing effects from averaging, we define the visual prototype for each class as the *medoid*—the instance whose embedding minimizes the average distance to others within the same class. For the prompt side, we collect the final trained verb embeddings for each class. Fig. 5 presents the average inter-class distance for both modalities across three conditions: (1) before training, (2) after training with CMMP [24], and (3) after training with our method (VDRP). We observe that CMMP significantly increases prompt diversity (0.937) while visual diversity remains moderate (0.420), indicating a modality mismatch. In contrast, VDRP maintains a comparable level of diversity in both modalities (prompt: 0.550, visual: 0.518), suggesting improved distributional alignment. This outcome implies that our prompt learning strategy—though primarily guided by intra-class visual variance—can indirectly induce a structured inter-class layout without over-amplifying semantic separation. We draw inspiration from recent works [60, 43, 36, 39, 40, 37], which emphasize distributional approaches to improving visual–prompt alignment. Motivated by this perspective, we interpret the improved visual–prompt alignment as a natural byproduct of our variance-aware design, even though cross-modal matching was not explicitly enforced.

Discussion. This inter-class alignment result further substantiates the novelty of our visual diversity-aware design. Unlike prior distribution-based prompt learning methods such as ProDA [37] and DAPT [36], which regularize textual embeddings, VDP explicitly grounds the prompt distribution in the visual modality by injecting group-wise variance extracted from union-region features. This design enables structured yet balanced prompt distributions that mirror the visual embedding space, achieving semantic coherence without explicit regularization losses. Consequently, VDP naturally promotes inter-class consistency and zero-shot generalization—key objectives that motivated our dual-module framework.

I Qualitative Results

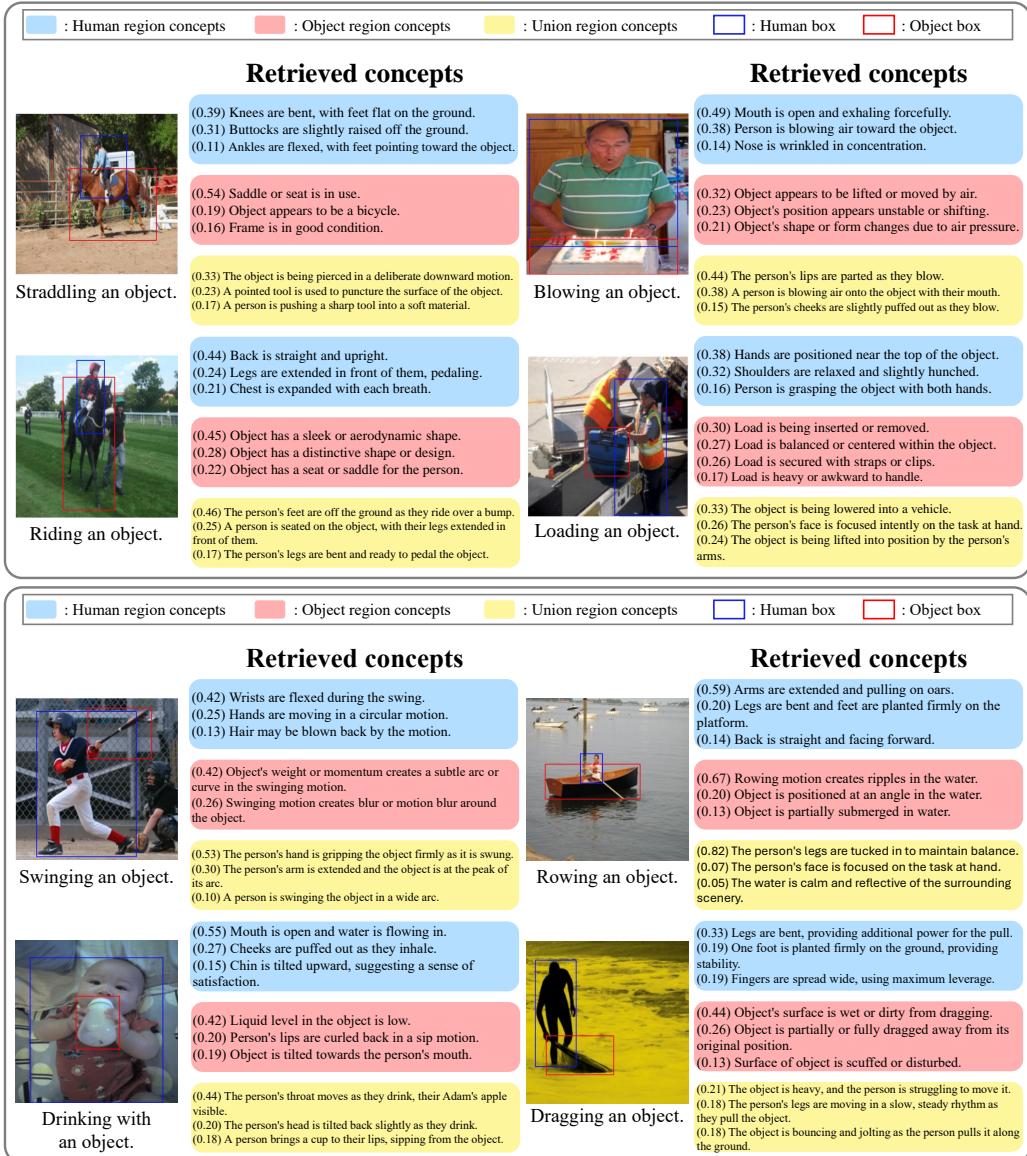


Figure 6: **Region-aware concept retrieval results (1/2).** Each row shows retrieved human, object, and union concepts. Concepts are color-coded by region: blue (human), red (object), yellow (union).

We present qualitative examples in Fig. 6 and 7 to illustrate how region-specific concepts contribute to verb classification. Each example visualizes the top-weighted concepts retrieved from three region branches—human, object, and union—based on sparsemax-normalized concept weights. The full



Figure 7: **Region-aware concept retrieval results (2/2).** Additional examples illustrating how different regions contribute to enhancing discriminability.

concept pool contains $K = 10$ entries per region type and verb class, and weights are displayed in parentheses. Due to sparsemax activation, concepts with low similarity scores are often assigned zero weight, resulting in more selective and interpretable retrieval. For instance, in *kicking an object*, the object region retrieves dynamic descriptions such as “(0.50) Object appears to be mid-air or suspended,” while the human region emphasizes posture concepts like “(0.27) Buttocks are tense and slightly lifted.” These complementary signals capture both motion and pose indicative of a kicking action. Similarly, in distinguishing *hugging* from *toasting*, human concepts like “(0.35) Arms wrap around the object” help identify close physical contact, whereas in *toasting*, object and union concepts highlight celebratory gestures such as “(0.90) Person's body language suggests celebration” and “(0.45) The person's smile is subtle as they toast the object.” These examples show how our model selectively retrieves fine-grained region-specific concepts that support verb disambiguation in subtle interaction contexts. Overall, these results show that our region-aware prompt augmentation (RAP) selects meaningful concepts from distinct regions, supporting both interpretability and fine-grained HOI discrimination.

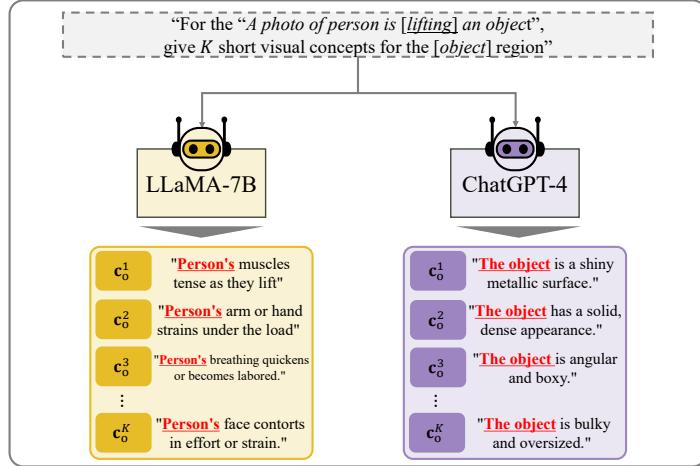


Figure 8: **Qualitative comparison of concept generation for the object region.** LLaMA-7B often yields human-centric concepts (left), while ChatGPT-4 provides more object-centered descriptions (right). This difference highlights the source of semantic noise in object-region concept retrieval.

Table 17: Comparison of performance when using object concepts from LLaMA-7B and ChatGPT-4 across four zero-shot settings. Best results are in **bold**.

NF-UC	Full	Unseen	Seen	RF-UC	Full	Unseen	Seen
LLaMA-7B	32.18	36.16	31.18	LLaMA-7B	33.55	31.24	34.13
ChatGPT-4	32.57	36.45	31.60	ChatGPT-4	33.78	31.29	34.41
UO	Full	Unseen	Seen	UV	Full	Unseen	Seen
LLaMA-7B	32.79	35.04	32.34	LLaMA-7B	32.78	26.88	33.74
ChatGPT-4	33.39	36.13	32.84	ChatGPT-4	32.73	26.69	33.72

J Analysis of object concept noise from LLaMA-7B

We further analyze the qualitative and quantitative effects of noisy object concepts generated by LLaMA-7B [48]. As discussed in the main paper, the prompt template “*a photo of a person [verb]-ing the object*” sometimes causes semantic leakage from the human region, producing human-centric descriptions for the object region. Fig. 8 visualizes this issue for the verb *lift*, where LLaMA-7B frequently associates object concepts with human body parts or actions. In contrast, ChatGPT-4 [49] produces more object-centered and visually grounded descriptions, effectively reducing such leakage. To quantitatively evaluate the impact of concept quality, we compare the model performance when object concepts are generated by LLaMA-7B versus ChatGPT-4. As shown in Table 17, GPT-4 generally achieves higher performance across all zero-shot settings, confirming that cleaner concept representations improve robustness without altering the model architecture.