

T-Rex-Omni: Integrating Negative Visual Prompt in Generic Object Detection

Jiazhou Zhou^{1,2†}, Qing Jiang^{2,3‡}, Kanghao Chen¹,
Lutao Jiang¹, Yuanhuiyi Lyu¹, Ying-Cong Chen¹, Lei Zhang^{2‡}

¹ Hong Kong University of Science and Technology (Guangzhou)

² International Digital Economy Academy (IDEA)

³ South China University of Technology

{zhoujiazhou, jiangqing, leizhang}@idea.edu.cn

Abstract

Object detection methods have evolved from closed-set to open-set paradigms over the years. Current open-set object detectors, however, remain constrained by their exclusive reliance on positive indicators based on given prompts like text descriptions or visual exemplars. This positive-only paradigm experiences consistent vulnerability to visually similar but semantically different distractors. We propose *T-Rex-Omni*, a novel framework that addresses this limitation by incorporating negative visual prompts to negate hard negative distractors. Specifically, we first introduce a unified visual prompt encoder that jointly processes positive and negative visual prompts. Next, a training-free Negating Negative Computing (NNC) module is proposed to dynamically suppress negative responses during the probability computing stage. To further boost performance through fine-tuning, our Negating Negative Hinge (NNH) loss enforces discriminative margins between positive and negative embeddings. *T-Rex-Omni* supports flexible deployment in both positive-only and joint positive-negative inference modes, accommodating either user-specified or automatically generated negative examples. Extensive experiments demonstrate remarkable zero-shot detection performance, significantly narrowing the performance gap between visual-prompted and text-prompted methods while showing particular strength in long-tailed scenarios (51.2 AP_r on LVIS-minival). This work establishes negative prompts as a crucial new dimension for advancing open-set visual recognition systems.

Extended version — <https://arxiv.org/html/2511.08997v1>

Introduction

Object detection stands as a cornerstone of computer vision, tasked with precisely localizing and categorizing objects within images. The field has evolved remarkably in recent years. Initially dominated by closed-set paradigms (Carion et al. 2020; Li et al. 2022b; Liu et al. 2022; Zhang et al. 2022; Zhu et al. 2020) limited to predefined categories, it has now shifted toward more flexible open-set detection systems (Li

et al. 2022c; Liu et al. 2024; Zhou et al. 2022) that can identify objects specified through user prompts, including text prompts (Gu et al. 2021; Li et al. 2022c; Liu et al. 2024; Zhan et al. 2024; Jiang et al. 2024b) (“a photo of a muffin”), visual prompts (Minderer et al. 2022; Xu et al. 2023; Zang et al. 2022; Li et al. 2024) (reference images of muffin), or combinations of them (Jiang et al. 2024a).

Modern open-set detection systems, however, face a fundamental limitation: their exclusive reliance on positive indicators from given prompts leaves them vulnerable to hard negatives—visually similar yet semantically distinct instances. As illustrated in Fig. 1, even state-of-the-art detectors may confidently classify a Chihuahua as a muffin when relying solely on positive visual prompts (e.g., a bound-

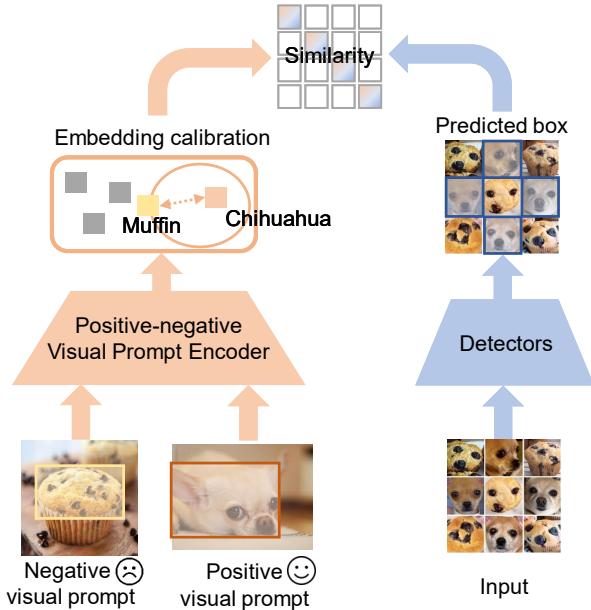


Figure 1: *T-Rex-Omni* employs dual visual prompts to enhance detection precision: positive prompts (e.g., “Chihuahua”) guide target object localization while negative prompts (e.g., “muffin”) actively suppress visually similar distractors. The joint positive-negative framework enables more discriminative and user-specified object detection.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[†]This work was done when Jiazhou Zhou and Qing Jiang were interns at IDEA.

[‡]Corresponding author.

ing box example of a Chihuahua). This issue is exacerbated by long-tailed data distributions, where detectors underperform for rare categories, leading to poor real-world applicability. To mitigate this, we propose leveraging negative visual prompts (e.g., a bounding box example of a muffin) to explicitly guide detectors away from hard negative distractors while preserving sensitivity to positive instances. This motivates our core research question: *Can visual negative prompts enable models to actively negate hard negatives without compromising their ability to detect true positives?*

Our solution, *T-Rex-Omni*, systematically integrates negative prompts into modern detection frameworks. First, we introduce a **unified positive-negative prompt encoder** that jointly processes positive and negative visual prompts from single or multiple images into corresponding prompt embeddings. To mitigate data scarcity, our prompt encoder synthesizes visual prompts by randomly jittering and resizing ground-truth boxes (mild for positives, strong for negatives). This augmentation enhances reference robustness to spatial variations and supports cross-image object detection. Next, we propose a **Negating Negative Computing (NNC) module**, which adaptively suppresses negative responses during probability computation. This training-free design permits immediate deployment with significant performance gains (Tab. 2). For further improvement via fine-tuning, a **Negating Negative Hinge (NNH) loss** enforces discriminative margins between positive and negative prompts in the embedding space, actively pushing apart visually similar but semantically distinct category embeddings (Tab. 2).

In this way, *T-Rex-Omni* offers three flexible visual prompt settings during inference: **(1) User-curated mode:** Users explicitly specify both positive and negative exemplars for precision-critical applications. **(2) Auto-suggested mode:** The system automatically proposes relevant negative exemplars based on user-provided positives, enabling efficient deployment with minimal user input. **(3) Positive-only mode:** Traditional single-prompt operation for rapid deployment. *T-Rex-Omni* tri-mode inference enables practitioners to dynamically adapt to varying precision and efficiency requirements across different application scenarios.

T-Rex-Omni exhibits strong object detection performance across four challenging benchmarks (COCO, LVIS, ODinW, Roboflow100) in zero-shot settings. Our key findings reveal that negative visual prompts can mitigate the previous modality gap between text-prompt and visual-prompt methods. *T-Rex-Omni* (Swim-L) even surpasses traditional text-prompt methods by **+2.0 AP** (LVIS-val) in Tab. 1. Besides, it delivers exceptional performance for rare categories on LVIS-minival (**51.2 AP_r**), LVIS-Val (**49.8 AP_r**), ODinW (**29.6 AP_{avg}**), and Roboflow (**20.3 AP_{avg}**), significantly outperforming existing approaches in long-tailed scenarios (Tab. 1). To summarize, our contributions are threefold:

- A simple yet effective framework that integrates negative examples as prompts for object detection, which can negate hard negative distractors while maintaining sensitivity to positive examples.
- An approach for embedding space modification through negative examples, supported by the plug-and-play NNC

module to suppress hard negative probabilities in a training-free manner and the NNH loss for enforced positive and negative embedding separation for fine-tuning adaptation.

- Extensive empirical validation across multiple benchmarks (COCO, LVIS, ODinW, Roboflow 100) demonstrating consistent zero-shot improvements, particularly in challenging long-tailed scenarios.

Related Work

Object Detection Object detection (OD) has transitioned from closed-set to open-set approaches to fit the dynamic and unpredictable nature of real-world environments. Early closed-set detection systems (Carion et al. 2020; Li et al. 2022b; Liu et al. 2022; Zhang et al. 2022; Zhu et al. 2020) are limited to recognizing objects from predefined categories. By contrast, open-set models adapt to identify objects beyond initial training categories. A prevalent method for open-set object detection leverages **text prompts** (Gu et al. 2021; Liu et al. 2024; Yao et al. 2022), which typically leverage knowledge from language models like CLIP (Radford et al. 2021) or BERT (Devlin et al. 2019) to align textual descriptions with visual representations. Recent multimodal large language models (MLLMs) (Hurst et al. 2024; Bai et al. 2025; Wu et al. 2024) have enhanced models with text referring expressions to capture object attributes, relationships, spatial configurations, and their interactions (Jiang et al. 2024b; Chen et al. 2023; Ma et al. 2024). However, texts often fail to accurately describe objects that are difficult to articulate. The **visual prompt-based** model provides a more intuitive representation through visual examples, such as image-level prompts (raw images) and object-level prompts (boxes, points). In this paper, we focus on the visual prompt-based object detection, which can benefit from open-vocabulary long-tailed objects.

Negative Sampling Negative sampling selects informative negatives or generates synthetic ones to improve representation learning while maintaining computational efficiency (Duan et al. 2024). This technique has demonstrated broad applicability across domains including recommendation systems (Yang et al. 2020a; Shi et al. 2023), natural language processing (Yang et al. 2024b; Zhan et al. 2021), graph learning (Yang et al. 2020c; Duan et al. 2022), and computer vision (Yang et al. 2024a; Wang et al. 2021). Negative sampling plays a pivotal role in object detection. Focal Loss (Lin et al. 2017) dynamically up-weights hard negatives during training. NP-RepMet (Yang et al. 2020b) jointly optimizes negative and positive prototypes for few-shot detection. UNP (Yan et al. 2024) isolates confusing negatives while ensuring the contribution of hard negatives via gradient modulation. GenNeg (Zhao et al. 2024a) leverages large-language models and text-to-image diffusion models to synthesize negative object descriptions and images. While these methods demonstrate the value of negative information for object detection, they are fundamentally constrained by their limited generalization capability for unseen categories due to their inherent dependence on training-time optimization. *T-Rex-Omni* introduces a fundamental shift by: (1)

enabling dynamic specification of negative samples using bounding box visual prompts; (2) supporting both training-free immediate deployment and fine-tunable versions for enhanced performance; and (3) excelling in long-tailed scenarios where previous methods struggle.

Model

Preliminary and Overview

Preliminary for *T-Rex2*. *T-Rex2* (Jiang et al. 2024a) adopts the DETR (Carion et al. 2020) framework, implementing an end-to-end Transformer-based architecture for open-set object detection. The model supports both joint and independent use of visual and text prompts through four key components: (1) a text prompt encoder based on CLIP’s text encoder, (2) a visual prompt encoder employing deformable cross-attention to process box/point interactions, (3) an image encoder for feature extraction, and (4) a DETR-style decoder for detection. *T-Rex2* enables generic and interactive visual prompt workflows for open-set object detection.

Overview for *T-Rex-Omni*. Built upon the *T-Rex2* architecture (Jiang et al. 2024a), *T-Rex-Omni* streamlines the framework by removing the text-prompt branch while introducing three novel components: (1) a unified prompt encoder for joint positive-negative prompt processing, (2) a training-free Negating Negative Computing (NNC) module for dynamic probability calibration, and (3) a Negating Negative Hinge (NNH) loss for discriminative embedding separation. As shown in Fig. 2, we preserve *T-Rex2*’s image encoder and DETR-style decoder, enhancing them with our proposed modules for negative prompt integration. The following subsections detail these contributions, with ablation studies in the experiment section validating their impacts.

Positive-negative Visual Prompt Encoder

T-Rex-Omni aims to learn both positive and negative visual prompt embeddings and integrate them into object detection. Building upon *T-Rex2*’s framework, we present several modifications to its original visual prompt encoder.

Visual Prompt Generation. During training, we generate both positive and negative visual prompts by performing random sampling and geometric jittering of the ground truth annotations. Specifically, for each category c present in an image, we first randomly sample one ground truth bounding box, denoted as $G_c = (x_c, y_c, w_c, h_c)$. Then, positive visual prompts p_c are synthesized by applying mild transformations, namely, random scaling or shifting within $[0, 0.3]$ scale range for the sampled ground truth bounding box G_c . Similarly, multiple negative visual prompts $n_c^i, i \in 1, 2, \dots, K$ are generated with stronger transformations within $[0.7, 1.0]$ scale range by repeating this process K times. This augmentation enriches prompt diversity while preserving semantic validity. Besides, it enhances robustness to spatial and scale variations in test-time visual prompts.

During inference, we maintain the user-provided positive visual prompt without modification. The negative visual prompts are obtained via consistent training augmentation (auto-suggested mode) or specified by users (user-curated mode).

Visual Prompt Encoder. The visual prompt encoder transforms both positive and negative visual prompts from coordinate space to embedding space. Given positive visual prompts p_c and K negative visual prompts $n_c^i, i \in 1, 2, \dots, K$, we initialize two learnable prompt queries $Q_P \in \mathbb{R}^{1 \times D}$ and $Q_N \in \mathbb{R}^{K \times D}$, where D is the dimension. Given the multi-scale image features $F = \{f_j | j \in 1, 2, \dots, L\}$ extracted from the image encoder with L feature map layers, we compute the enhanced queries via multi-scale deformable cross-attention (Zhu et al. 2020). The positive prompt queries Q'_P and negative prompt queries Q'_N are computed as:

$$Q'_P = MSDeformAttn(Q_P, p_c, F), \quad (1)$$

$$Q'_N = MSDeformAttn(Q_N, n_c, F), \quad (2)$$

The deformable attention allows each prompt query to dynamically attend to the most relevant image features within its respective visual prompt region. These attended features are then refined through a self-attention layer followed by a feed-forward network (FFN). The positive V_P and negative V_N prompt embeddings are obtained as follows:

$$V_P = FFN(SelfAttn(Q'_P)), \quad (3)$$

$$V_N = FFN(SelfAttn(Q'_N)), \quad (4)$$

The above operations are repeated across all categories for each image. Given a training batch of B images with a maximum of M categories per batch, we denote the positive prompts as $V'_P \in \mathbb{R}^{M \times B \times D}$ and K negative visual prompts as $V'_N \in \mathbb{R}^{M \times B \times K \times D}$.

To facilitate *T-Rex-Omni* for cross-image object detection, we enforce that each training batch contains at least one shared category (see Sec. 4.1). This enables us to compute positive averaged prompt embeddings of shared categories across the batch, namely $V''_P \in \mathbb{R}^{M \times D}$, thus effectively propagating prompt embeddings between images of the same category. For negative prompt embeddings, we select the top- K most similar embeddings to the averaged positive embedding V''_P from the full set of $B \times K$ candidate negatives, yielding the final negative embeddings $V''_N \in \mathbb{R}^{M \times K \times D}$. The above batch-wise operation is disabled during inference. Overall, the V''_P captures the target object’s visual characteristics while V''_N represents visually similar but incorrect or suboptimal detections that should be distinguished from the target.

Flexible Inference. During deployment, *T-Rex-Omni* supports three visual prompt settings to accommodate diverse application requirements: (1) **User-curated mode:** Users explicitly specify both positive and negative exemplars through bounding box annotations or precise image crops. This high-precision mode is particularly valuable for critical applications. (2) **Auto-suggested mode:** The system automatically proposes relevant negative exemplars by applying geometric transformations to user-provided positive exemplars. This balanced mode reduces annotation effort and serves as our default setting for benchmark evaluations.

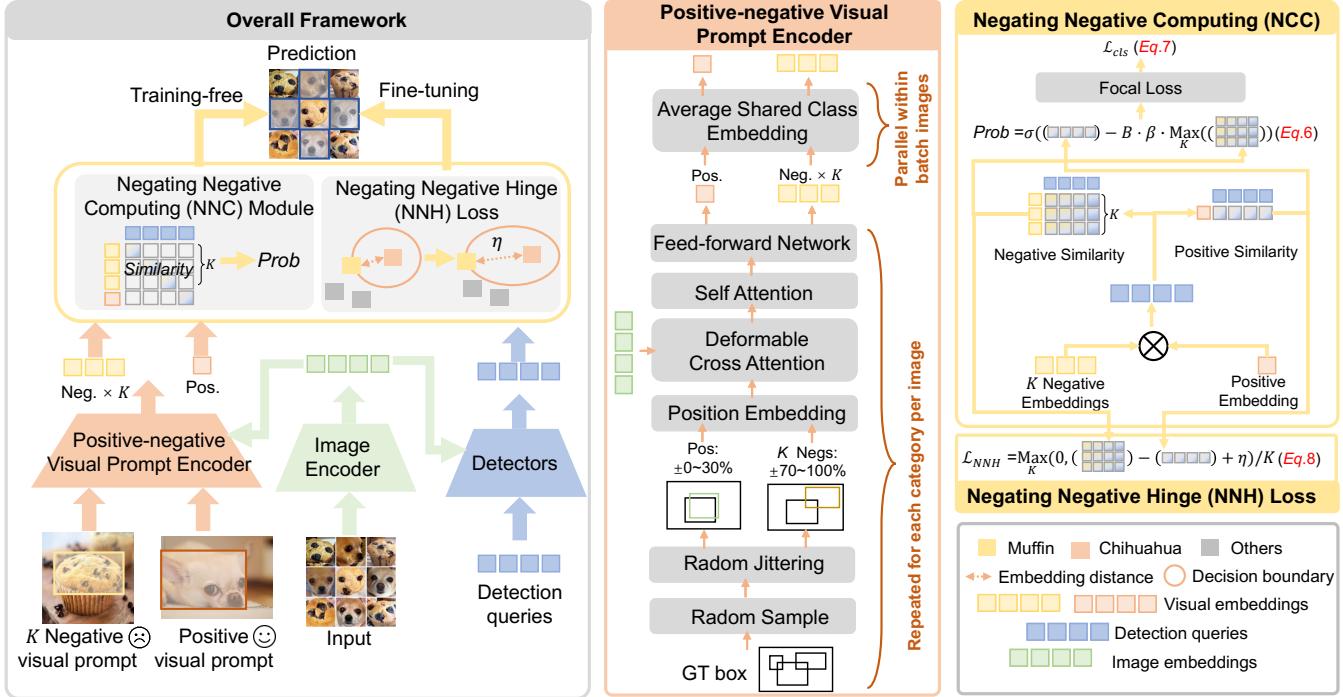


Figure 2: Overview of the *T-Rex-Omni* model.

(3) Positive-only mode: Maintains compatibility with traditional positive-only prompt workflows for rapid deployment scenarios. *T-Rex-Omni* tri-mode inference enables practitioners to dynamically adapt to varying precision and efficiency requirements across different application scenarios.

Negating Negative Computing Module.

Built upon the extracted positive V_P'' and K negative prompt embeddings $V_{N,i}''$, $i = 1, 2, \dots, K$, our NNC module calibrates prediction confidence by suppressing scores for hard negative cases. As shown in Fig. 2, given N_q detection queries $Q \in \mathbb{R}^{N_q \times D_q}$ from DETR decoder, we compute positive S_P and negative $S_{N,i}$ similarity scores between detection queries with both positive and multiple negative prompt embeddings via matrix multiplication:

$$S_P, S_{N,i} = Q \times (V_P'')^T, Q \times (V_{N,i}'')^T \quad (5)$$

We then subtract the positive similarity scores from the weighted negative ones to suppress the likelihood of predicting classes that are visually similar but semantically different. The following sigmoid function σ then transforms the similarity score into probability $Prob$ as follows:

$$Prob = \sigma(S_P - B \cdot \beta \cdot \max_{i=1,2,\dots,K}(S_{N,i})), \quad (6)$$

where $0 < \beta < 1$ is the introduced parameter to control how much the negative examples influence the final score; the max operation $\max(\cdot)$ selects the strongest negative similarity across K negative similarities and $B \sim \text{Bernoulli}(0.5)$ is a stochastic indicator for mode switching. During training, B stochastically switches between joint

positive-negative ($B = 1$) and positive-only ($B = 0$) modes to ensure inference compatibility.

In training-free or inference applications, NNC operates as a plug-and-play module, demonstrating consistent performance without fine-tuning (Tab. 2). For fine-tuning training applications, we integrate the predicted probabilities from the NNC module into Focal loss (Lin et al. 2017) to calculate the classification loss and backpropagation:

$$\mathcal{L}_{cls} = -\alpha_t(1 - Prob_t)^\gamma \log(Prob_t), \quad (7)$$

where $Prob_t$ is defined as $Prob_t = Prob$ if the class prediction is true, otherwise $Prob_t = 1 - Prob$; $\alpha_t \in [0, 1]$ is the class-balancing weight and $\gamma \geq 0$ controls the focus on hard examples.

Negating Negative Hinge Loss.

To improve discrimination between visually similar but semantically distinct categories, we propose the Negating Negative Hinge (NNH) loss, which explicitly enforces a margin-based separation between positive and negative prompt embeddings. The loss is defined as:

$$\mathcal{L}_{Hinge} = \sum_{i=1,2,\dots,K} \text{Max}(0, S_{N,i} - S_P + \eta)/K \quad (8)$$

where S_P and $S_{N,i}$ are the similarity scores for the positive and the i -th negative prompt embedding; $\eta > 0$ is a pre-set margin that controls the minimum separation between positive and negative similarities; and K is the number of negative prompt embeddings.

The NNH loss ensures that the similarity S_P for the positive embeddings exceeds the similarity $S_{N,i}$ for any negative

embeddings by at least the margin η . This constraint encourages the model to learn more discriminative embeddings by penalizing cases where negative similarities overlap with positive similarities. The hinge loss formulation $\text{Max}(0, \cdot)$ provides a robust optimization objective, as it only penalizes violations of the margin condition, making the training process more stable and focusing on hard negative cases.

Training Strategy and Objective.

Training Strategy. Unlike *T-Rex2* (Jiang et al. 2024a), which employs a “current image prompt, current image detect” training paradigm, we introduce a “current image prompt, cross-image detect” training strategy. By ensuring each training batch contains at least one shared category across images (see Sec. 4.1), our approach encourages more robust visual prompt learning through inter-image consistency, thus enhancing cross-image object detection capability and generalization of prompt embeddings.

Training Objectives. Our complete loss function combines box regression losses (L1 and GIoU (Rezatofighi et al. 2019)), classification loss \mathcal{L}_{cls} (Eq. 9), our NNC loss \mathcal{L}_{Hinge} (Eq. 10), auxiliary losses (intermediate supervision after each decoder layer and encoder outputs), and denoising training loss proposed in DINO (Zhang et al. 2022) to accelerate convergence. The box regression and classification loss are initially employed for bipartite matching (Carion et al. 2020) between predictions and ground truths. The final objective function is:

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \mathcal{L}_{Hinge} + \mathcal{L}_{L1} + \mathcal{L}_{GIoU} + \mathcal{L}_{DN} \quad (9)$$

Experiments

Data Engine

Our visual prompt object detection framework employs a specialized batch construction strategy to enable cross-image detection. Each training batch contains images sharing at least one object category, allowing object instances from one image to serve as visual prompts for detecting corresponding instances in other batch images. The batch construction involves two key steps: (1) building a hash table that maps object categories to images containing more than three instances to ensure instance diversity, and (2) for each image, selecting its second-most frequent category and retrieving matching images from the hash table. We apply this strategy to generate training batches for fine-tuning on the Objects365 dataset (Shao et al. 2019).

Model Details

T-Rex-Omni adopts the pre-trained weights from *T-Rex2* (Jiang et al. 2024a), featuring a Swin Transformer (Liu et al. 2021) backbone and six-layer Transformer encoder. Designed specifically for visual-prompt object detection, our architecture omits text encoders and employs: (1) a visual prompt encoder with three deformable cross-attention layers with a hidden dimension set to 1024 and (2) a prompt sampling strategy that selects one ground-truth instance box as a positive prompt and three randomly jittered boxes as negative prompts per category. The hyperparameters β for

the NNC module, η for the NNC loss, α_t and γ for the focal loss are 0.3, 0.3, 0.25, and 2, respectively. We optimize using AdamW (Loshchilov and Hutter 2017) with differential learning rates (10^{-5} backbone, 10^{-4} others).

Settings and Metrics

We conduct zero-shot evaluation where the training images used by *T-Rex-Omni* has no overlap with the evaluation dataset. We report Average Precision (AP) metrics across four benchmarks: COCO (Lin et al. 2014), LVIS (Gupta, Dollar, and Girshick 2019), ODinW35 (Li et al. 2022a), and Roboflow100 (Ciaglia et al. 2022). We employ the following visual prompt settings for evaluation:

Visual-G: In this setting, we adhere to *T-Rex2* (Jiang et al. 2024a) Visual-G evaluation protocol for open-set object detection. For each benchmark category, we extract both positive and negative visual prompt embeddings from the training set images, with negatives created by randomly jittering ground truth boxes. Taking COCO as a representative example, this visual prompt embedding generation process follows three steps: (1) **Sampling:** For each category, we randomly select $N = 16$ images containing at least one instance of that category from the training dataset. (2) **Embedding Extraction:** We extract positive and three negative visual embeddings using each image’s ground truth box and corresponding jittered boxes as input, respectively. (3) **Aggregation:** We compute category-level average embeddings, thus yielding 80 positive embeddings and $80 \times K$ negative embeddings for COCO, where we set $K = 3$ in this evaluation based on our ablation study (Fig. 3d). This process is performed once and remains fixed during evaluation, ensuring consistent prompt representation across the benchmark.

Main Results

Zero-Shot Generic Object Detection. We conduct a comprehensive evaluation of *T-Rex-Omni*’s zero-shot capabilities (denoted as “Visual-G” in Tab. 1) across four challenging benchmarks, where “zero-shot” denotes evaluation on images excluded from training. Our analysis reveals four key findings: (1) **State-of-the-Art Visual Prompting:** With Swin-T, *T-Rex-Omni* outperforms the previous best visual-prompt approach (*T-Rex2*) by **+4.8 AP** (43.6 vs 38.8) on COCO-val and **+5.6 AP** (43.0 vs 37.4) on LVIS-minival, setting new visual-prompt benchmarks. (2) **Long-Tailed Superiority:** The most striking improvement appears in long-tailed scenarios. *T-Rex-Omni* achieves **+7.1 AP_r** for LVIS-minival rare categories (37.0 vs 29.9), demonstrating 23.8% relative improvement in long-tailed scenarios. (3) **Text-Visual Gap Reduction:** *T-Rex-Omni* narrows the text-visual performance gap to **2.2 AP** on COCO (43.6 vs *T-Rex2*-text’s 45.8) while surpassing text-prompt *T-Rex2* on LVIS-val by **+2.8 AP** (37.7 vs 34.9). (4) **Backbone Scalability:** With the Swin-L backbone, improvements remain consistent: **+4.2 AP** on COCO (50.7 vs 46.5) and **+5.8 AP_r** for LVIS rare categories (51.2 vs 45.4), with ODinW-35 gains of **+1.8 AP_{avg}**.

Ablation Experiments

Ablation Study on the NNC Module and NNH Loss. As shown in Tab. 2, our NNC module significantly improves

Methods	Prompt Type	COCO-Val	LVIS								ODinW35	Roboflow100	
		Zero-Shot	Zero-Shot								Zero-Shot	Zero-Shot	
		val-80	minival-804				val-1203				val-35	val-100	
		AP	AP	AP _f	AP _c	AP _r	AP	AP _f	AP _c	AP _r	AP _{avg}	AP _{med}	
Swin-T Backbone													
GLIP-T (Li et al. 2022c)	Text	46.7	26.0	31.0	21.4	20.8	17.2	25.5	12.5	10.1	19.6	5.1	
Grounding DINO (Liu et al. 2023)	Text	48.4	27.4	32.7	23.3	18.1	-	-	-	-	22.3	11.9	
DetCLIPv2 (Yao et al. 2023)	Text	-	40.4	40.0	41.7	36.0	-	-	-	-	-	-	
MM-GDINO (Zhao et al. 2024b)	Text	-	41.4	46.2	37.4	34.2	31.9	40.5	27.6	23.6	23.1	-	
T-Rex2 (Jiang et al. 2024a)	Text	45.8	42.8	46.5	39.7	37.4	34.8	41.2	31.5	29.0	18.0	4.7	
LLMDet (Fu et al. 2025)	Text	-	44.7	50.7	39.5	37.3	34.9	44.3	30.1	26.0	23.8	-	
DINOv (Li et al. 2023)	Visual-G	-	-	-	-	-	-	-	-	-	14.9	5.4	
T-Rex2 (Jiang et al. 2024a)	Visual-G	38.8	37.4	41.8	33.9	29.9	34.9	41.1	30.3	32.4	23.6	17.5	
VisTex-DINO (Wu et al. 2025)	Text+Visual-G	-	42.8	-	-	37.2	-	-	-	-	-	-	
<i>T-Rex-Omni</i>	Visual-G	43.6	43.0	47.7	38.9	37.0	37.7	41.9	33.6	38.6	25.2	20.1	
Swin-L Backbone													
GLIP-L (Li et al. 2022c)	Text	49.8	37.3	41.5	34.3	28.2	26.9	35.4	23.3	17.1	23.4	11.0	
Grounding DINO (Liu et al. 2023)	Text	52.5	33.9	38.8	30.7	22.2	-	-	-	-	26.1	18.4	
DetCLIPv2 (Yao et al. 2023)	Text	-	44.7	43.7	46.3	43.1	-	-	-	-	-	-	
MM-GDINO (Zhao et al. 2024b)	Text	-	36.8	42.8	31.8	28.1	29.1	37.2	25.6	19.7	-	-	
LLMDet (Fu et al. 2025)	Text	-	51.1	56.6	46.1	45.1	42.0	50.2	38.8	31.6	-	-	
T-Rex2 (Jiang et al. 2024a)	Text	52.2	54.9	56.1	54.8	49.2	45.8	50.2	43.2	42.7	22.0	7.3	
DINOv (Li et al. 2023)	Visual-G	-	-	-	-	-	-	-	-	-	15.7	4.8	
T-Rex2 (Jiang et al. 2024a)	Visual-G	46.5	47.6	49.5	46.0	45.4	45.3	49.5	42.0	43.8	27.8	20.5	
VisTex-GLIP (Wu et al. 2025)	Text+Visual-G	-	50.7	-	-	42.9	-	-	-	-	-	-	
<i>T-Rex-Omni</i>	Visual-G	50.7	54.0	56.0	52.4	51.2	47.8	50.0	45.1	45.1	29.6	23.1	
20.3													

Table 1: Zero-shot object detection results. Best visual-prompted results are marked in **bold**.

Settings	COCO-val	LVIS-minival				
	Zero-Shot	Zero-Shot	AP	AP _r	AP _c	AP _f
NNC	✗	✗	38.8	37.4	41.8	33.9
NNH	✗	✗	41.8 ±3.0	40.6 ±3.2	45.2 ±3.4	36.9 ±3.0
Fine-tune	✓	✗	42.9 ±4.1	41.4 ±4.0	46.2 ±4.4	38.1 ±4.2
	✓	✓	43.6 ±4.8	43.0 ±5.6	47.7 ±5.9	38.9 ±5.0
						37.0 ±7.1

Table 2: Ablation study on the NNC module and NNH loss in fine-tuning and training-free modes.

performance in a training-free setting (row 2), increasing COCO-val AP by **+3.0** and LVIS-minival AP by **+3.2**. This validates NNC’s probability calibration (Eq. 6) in suppressing hard negative predictions. Further fine-tuning (row 3 using Eq. 9) enhances performance (COCO-val: **+1.1**, LVIS-minival: **+0.8**) while the full model (row 4), which incorporates the NNH loss (Eq. 8), achieves the best results (COCO-val: **43.6** AP, LVIS-minival: **43.0** AP), demonstrating the NNH loss’s ability to enforce discriminative margins in embedding space (Eq. 8). These results validate NNC’s stochastic negative suppression and fine-tuning jointly reduce false positives, while NNH’s margin separation enhances inter-class positive-negative discrimination.

Ablation Study on Prompt Setting Compatibility. Our systematic evaluation (Tab. 3) reveals three key findings regarding the mode switching training mechanism. First, models trained with fixed negative prompt integration ($B = 1$) show better performance in positive-negative evaluation (**42.4** AP on COCO-val) than in positive-only mode (**40.8** AP on COCO-val $\Delta = +1.6$), confirming the value of negative prompt utilization. Second, stochastic training yields

Settings		COCO-val	LVIS-minival				
		Zero-Shot	Zero-Shot	AP	AP _r	AP _c	AP _f
Training	Evaluation	AP	AP	AP _r	AP _c	AP _f	
$B = 0$	$B = 0$	38.8	37.4	41.8	33.9	29.9	
	$B = 1$	41.8	40.6	45.2	36.9	33.3	
$B = 1$	$B = 0$	40.8	39.8	44.0	35.7	31.1	
	$B = 1$	42.4	40.9	45.5	37.5	34.2	
$B(0.5)$	$B = 0$	42.0	41.0	45.9	37.2	34.8	
	$B = 1$	43.6	43.0	47.7	38.9	37.0	

Table 3: Ablation study on *T-Rex-Omni*’s compatibility for positive-only ($B = 0$) and positive-negative ($B = 1$) visual prompt settings. $B \sim \text{Bernoulli}(0.5)$ is an indicator for stochastic mode switching.

more robust performance across both evaluation settings, achieving better results in the positive-negative configuration (**43.6** AP on COCO-val, **+2.8** over fixed training). Most notably, the approach demonstrates exceptional performance on rare categories of LVIS-minival (**47.7** AP_r), with a **6.9** point improvement over fixed training, highlighting its effectiveness for long-tailed recognition. These consistent gains across COCO-val and LVIS-minival demonstrate that our mode switching training mechanism successfully bridges the gap between different prompt configurations while maintaining superior discriminative capabilities. **Hyperparameter β Selection in the NNC module.** We ablate the negative suppression coefficient β in the NNC module (Fig. 3a). Results show a unimodal performance curve, peaking at $\beta = 0.3$ (42.8 AP on COCO-val, 43.0 AP on LVIS-minival), balancing positive reinforcement and negative suppression. Performance degrades at extremes: no suppression ($\beta = 0.0$) yields 39.7 AP, while over-suppression

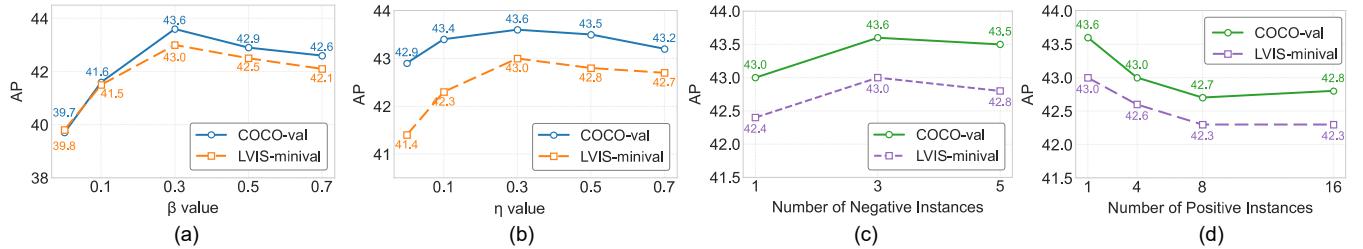


Figure 3: Ablation study on hyperparameters. (a) β in the NNC module; (b) η in the NNH loss; (c) positive prompt quantity; (d) negative prompt quantity.

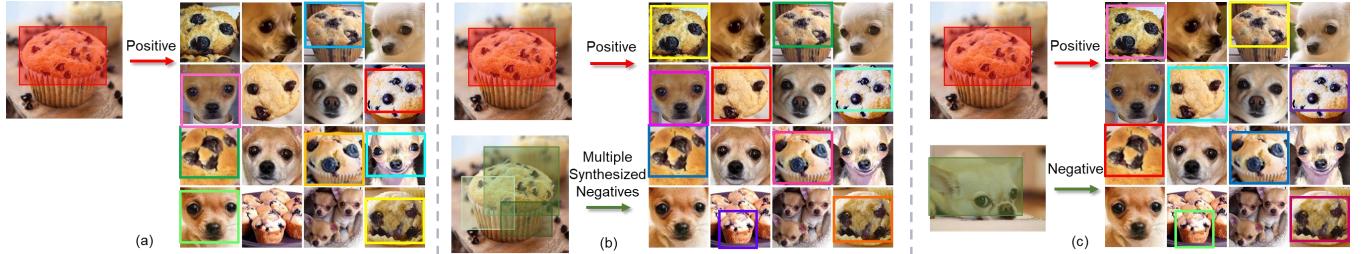


Figure 4: Visualization of *T-Rex-Omni*’s three inference modes. (a) Positive-only; (b) Auto-suggested; (c) User-curated.

($\beta \geq 0.5$) causes decline. This confirms that moderate suppression ($\beta \approx 0.3$) is critical—both neglecting and over-penalizing negatives harm discriminative power.

Hyperparameter η Selection in the NNH loss. As shown in Fig. 3b, the optimal detection result peaks at $\eta = 0.3$ (43.6 AP on COCO-val, 43.0 AP on LVIS-minival). A moderate margin ($\eta = 0.1 - 0.3$) outperforms the no margin setting ($\eta = 0$), confirming that our enforced embedding separation enhances discriminative power. However, excessive margins ($\eta \geq 0.5$) degrade performance, suggesting overly aggressive separation harms embedding representations. The peak at $\eta = 0.3$ balances discriminative features and semantic relationships, validating our choice of η in the NNH loss.

Impact of Negative Example Quantity on Model Performance. As shown in Fig. 3c, our ablation study shows that employing three negative prompts yields optimal zero-shot detection accuracy (43.6 AP on COCO-val, 43.0 AP on LVIS-minival), representing a 0.6 AP improvement over the single-negative baseline. However, further increasing to five negative examples yields diminishing returns (43.5 AP on COCO-val), suggesting an upper bound on the benefits of negative prompt diversity. This pattern indicates that while multiple negative examples help discriminate between visually similar categories, excessive negative prompts may introduce noise or redundant information.

Impact of Positive Prompt Quantity on Model Performance. Our ablation study reveals a counterintuitive relationship between the number of positive visual prompts and zero-shot detection performance. As shown in Fig. 3d, the results indicate that using a single positive example yields optimal performance (43.6 AP on COCO-val, 43.0 AP on LVIS-minival), with progressively degraded results as more examples are incorporated. Specifically, increasing the number of positive examples to 4, 8, and 16 leads to performance

drops of 0.6, 0.9, and 0.8 AP points, respectively, on COCO-val, with similar degradation patterns observed on LVIS-minival. This suggests that *T-Rex-Omni* benefits more from high-quality positive examples rather than quantity, potentially due to reduced noise in the learned representations.

Visualization of *T-Rex-Omni*’s three inference modes. Fig. 4a demonstrates *T-Rex-Omni*’s positive-only baseline (using muffin prompts) yields imperfect results, missing some muffins and misclassifying chihuahuas. With Auto-suggested mode (Fig. 4b), *T-Rex-Omni* automatically synthesizes multiple negative visual prompts conditioned on the positive prompt, generating refined detections with fewer errors. Most impressively, for user-curated mode (Fig. 4c), *T-Rex-Omni* leverages both user-provided positive and negative prompts to achieve optimal performance—correctly detecting all muffins while eliminating all chihuahua misclassifications. This shows *T-Rex-Omni*’s flexible visual prompt settings and progressive improvement.

Conclusion

We have introduced *T-Rex-Omni*, a novel framework that advances open-set object detection through the integration of negative visual prompts. Addressing a critical limitation of existing positive-only paradigms—their vulnerability to visually similar distractors, *T-Rex-Omni* achieves enhanced detection performance via three key contributions: (1) a training-free NNC module for hard negative suppression, (2) an NNH loss for embedding space regularization, and (3) a unified architecture for joint positive-negative prompt processing. Extensive experiments demonstrate significantly enhanced robustness, evidenced by both a reduced performance gap between visual and text prompts and superior performance in long-tailed scenarios.

Acknowledgement

This work is partially supported by the National Natural Science Foundation of China (No. 62206068).

References

Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229. Springer.

Chen, K.; Zhang, Z.; Zeng, W.; Zhang, R.; Zhu, F.; and Zhao, R. 2023. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*.

Ciaglia, F.; Zuppichini, F. S.; Guerrie, P.; McQuade, M.; and Solaetz, J. 2022. Roboflow 100: A Rich, Multi-Domain Object Detection Benchmark. *arXiv preprint arXiv:2211.13523*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.

Duan, S.; Yi, X.; Zhang, P.; Liu, Y.; Liu, Z.; Lu, T.; Xie, X.; and Gu, N. 2024. Negating Negatives: Alignment with Human Negative Samples via Distributional Dispreference Optimization. *arXiv preprint arXiv:2403.03419*.

Duan, W.; Xuan, J.; Qiao, M.; and Lu, J. 2022. Learning from the dark: boosting graph convolutional neural networks with diverse negative samples. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 6550–6558.

Fu, S.; Yang, Q.; Mo, Q.; Yan, J.; Wei, X.; Meng, J.; Xie, X.; and Zheng, W.-S. 2025. Llmdet: Learning strong open-vocabulary object detectors under the supervision of large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 14987–14997.

Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2021. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*.

Gupta, A.; Dollar, P.; and Girshick, R. 2019. LVIS: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5356–5364.

Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.

Jiang, Q.; Li, F.; Zeng, Z.; Ren, T.; Liu, S.; and Zhang, L. 2024a. T-rex2: Towards generic object detection via text-visual prompt synergy. In *European Conference on Computer Vision*, 38–57. Springer.

Jiang, Q.; Luo, G.; Yang, Y.; Xiong, Y.; Chen, Y.; Zeng, Z.; Ren, T.; and Zhang, L. 2024b. Chatrex: Taming multimodal llm for joint perception and understanding. *arXiv preprint arXiv:2411.18363*.

Li, C.; Liu, H.; Li, L. H.; Zhang, P.; Aneja, J.; Yang, J.; Jin, P.; Lee, Y. J.; Hu, H.; Liu, Z.; et al. 2022a. ELEVATER: A Benchmark and Toolkit for Evaluating Language-Augmented Visual Models. *arXiv preprint arXiv:2204.08790*.

Li, F.; Jiang, Q.; Zhang, H.; Ren, T.; Liu, S.; Zou, X.; Xu, H.; Li, H.; Li, C.; Yang, J.; et al. 2023. Visual In-Context Prompting. *arXiv preprint arXiv:2311.13601*.

Li, F.; Jiang, Q.; Zhang, H.; Ren, T.; Liu, S.; Zou, X.; Xu, H.; Li, H.; Yang, J.; Li, C.; et al. 2024. Visual in-context prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12861–12871.

Li, F.; Zhang, H.; Liu, S.; Guo, J.; Ni, L. M.; and Zhang, L. 2022b. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13619–13627.

Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022c. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10965–10975.

Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.

Liu, S.; Li, F.; Zhang, H.; Yang, X.; Qi, X.; Su, H.; Zhu, J.; and Zhang, L. 2022. DAB-DETR: Dynamic Anchor Boxes are Better Queries for DETR. In *International Conference on Learning Representations*.

Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, 38–55. Springer.

Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.

Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.

Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.

Ma, C.; Jiang, Y.; Wu, J.; Yuan, Z.; and Qi, X. 2024. Groma: Localized visual tokenization for grounding multi-modal large language models. In *European Conference on Computer Vision*, 417–435. Springer.

Minderer, M.; Gritsenko, A.; Stone, A.; Neumann, M.; Weissenborn, D.; Dosovitskiy, A.; Mahendran, A.; Arnab, A.; Dehghani, M.; Shen, Z.; et al. 2022. Simple open-vocabulary object detection. In *European conference on computer vision*, 728–755. Springer.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.

Rezatofighi, H.; Tsai, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 658–666.

Shao, S.; Li, Z.; Zhang, T.; Peng, C.; Yu, G.; Zhang, X.; Li, J.; and Sun, J. 2019. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8430–8439.

Shi, W.; Chen, J.; Feng, F.; Zhang, J.; Wu, J.; Gao, C.; and He, X. 2023. On the theories behind hard negative sampling for recommendation. In *Proceedings of the ACM Web Conference 2023*, 812–822.

Wang, W.; Zhou, T.; Yu, F.; Dai, J.; Konukoglu, E.; and Van Gool, L. 2021. Exploring cross-image pixel contrast for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7303–7313.

Wu, Y.; Zhou, Y.; Saiyin, J.; Wei, B.; and Xu, Y. 2025. Visual Textualization for Image Prompted Object Detection. *arXiv preprint arXiv:2506.23785*.

Wu, Z.; Chen, X.; Pan, Z.; Liu, X.; Liu, W.; Dai, D.; Gao, H.; Ma, Y.; Wu, C.; Wang, B.; et al. 2024. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*.

Xu, Y.; Zhang, M.; Fu, C.; Chen, P.; Yang, X.; Li, K.; and Xu, C. 2023. Multi-modal queried object detection in the wild. *Advances in Neural Information Processing Systems*, 36: 4452–4469.

Yan, B.; Lang, C.; Cheng, G.; and Han, J. 2024. Understanding negative proposals in generic few-shot object detection. *IEEE Transactions on Circuits and Systems for Video Technology*.

Yang, J.; Yi, X.; Zhiyuan Cheng, D.; Hong, L.; Li, Y.; Xiaoming Wang, S.; Xu, T.; and Chi, E. H. 2020a. Mixed negative sampling for learning two-tower neural networks in recommendations. In *Companion proceedings of the web conference 2020*, 441–447.

Yang, Y.; Wei, F.; Shi, M.; and Li, G. 2020b. Restoring negative information in few-shot object detection. *Advances in neural information processing systems*, 33: 3521–3532.

Yang, Z.; Ding, M.; Huang, T.; Cen, Y.; Song, J.; Xu, B.; Dong, Y.; and Tang, J. 2024a. Does negative sampling matter? a review with insights into its theory and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Yang, Z.; Ding, M.; Zhou, C.; Yang, H.; Zhou, J.; and Tang, J. 2020c. Understanding negative sampling in graph representation learning. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 1666–1676.

Yang, Z.; Shao, Z.; Dong, Y.; and Tang, J. 2024b. TriSampler: a better negative sampling principle for dense retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9269–9277.

Yao, L.; Han, J.; Liang, X.; Xu, D.; Zhang, W.; Li, Z.; and Xu, H. 2023. Detclipv2: Scalable open-vocabulary object detection pre-training via word-region alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23497–23506.

Yao, L.; Han, J.; Wen, Y.; Liang, X.; Xu, D.; Zhang, W.; Li, Z.; Xu, C.; and Xu, H. 2022. DetCLIP: Dictionary-Enriched Visual-Concept Paralleled Pre-training for Open-world Detection. *arXiv preprint arXiv:2209.09407*.

Zang, Y.; Li, W.; Zhou, K.; Huang, C.; and Loy, C. C. 2022. Open-vocabulary detr with conditional matching. In *European Conference on Computer Vision*, 106–122. Springer.

Zhan, J.; Mao, J.; Liu, Y.; Guo, J.; Zhang, M.; and Ma, S. 2021. Optimizing dense retrieval model training with hard negatives. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 1503–1512.

Zhan, Y.; Zhu, Y.; Chen, Z.; Yang, F.; Tang, M.; and Wang, J. 2024. Griffon: Spelling out all object locations at any granularity with large language models. In *European Conference on Computer Vision*, 405–422. Springer.

Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H.-Y. 2022. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. *arXiv:2203.03605*.

Zhao, S.; Zhao, L.; Suh, Y.; Metaxas, D. N.; Chandraker, M.; Schulter, S.; et al. 2024a. Generating enhanced negatives for training language-based object detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13592–13602.

Zhao, X.; Chen, Y.; Xu, S.; Li, X.; Wang, X.; Li, Y.; and Huang, H. 2024b. An open and comprehensive pipeline for unified object grounding and detection. *arXiv preprint arXiv:2401.02361*.

Zhou, X.; Girdhar, R.; Joulin, A.; Krähenbühl, P.; and Misra, I. 2022. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, 350–368. Springer.

Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.