

# LaMI-DETR: Open-Vocabulary Detection with Language Model Instruction

Penghui Du<sup>1,2,3\*</sup>, Yu Wang<sup>2\*</sup>, Yifan Sun<sup>2</sup>, Luting Wang<sup>1</sup>, Yue Liao<sup>1</sup>, Gang Zhang<sup>2</sup>, Errui Ding<sup>2</sup>, Yan Wang<sup>3</sup><sup>✉</sup>, Jingdong Wang<sup>2</sup>, and Si Liu<sup>1</sup><sup>✉</sup>

<sup>1</sup>Beihang University <sup>2</sup>Baidu <sup>3</sup>AIR, Tsinghua University  
{dupenghui, wangluting, liusi}@buaa.edu.cn liaoyue.ai@gmail.com  
{wangyu106, sunyifan01, zhanggang03, dingerrui, wangjingdong}@baidu.com  
wangyan@air.tsinghua.edu.cn

**Abstract.** Existing methods enhance open-vocabulary object detection by leveraging the robust open-vocabulary recognition capabilities of Vision-Language Models (VLMs), such as CLIP. However, two main challenges emerge: (1) A deficiency in concept representation, where the category names in CLIP’s text space lack textual and visual knowledge. (2) An overfitting tendency towards base categories, with the open vocabulary knowledge biased towards base categories during the transfer from VLMs to detectors. To address these challenges, we propose the Language Model Instruction (LaMI) strategy, which leverages the relationships between visual concepts and applies them within a simple yet effective DETR-like detector, termed LaMI-DETR. LaMI utilizes GPT to construct visual concepts and employs T5 to investigate visual similarities across categories. These inter-category relationships refine concept representation and avoid overfitting to base categories. Comprehensive experiments validate our approach’s superior performance over existing methods in the same rigorous setting without reliance on external training resources. LaMI-DETR achieves a rare box AP of 43.4 on OV-LVIS, surpassing the previous best by 7.8 rare box AP.

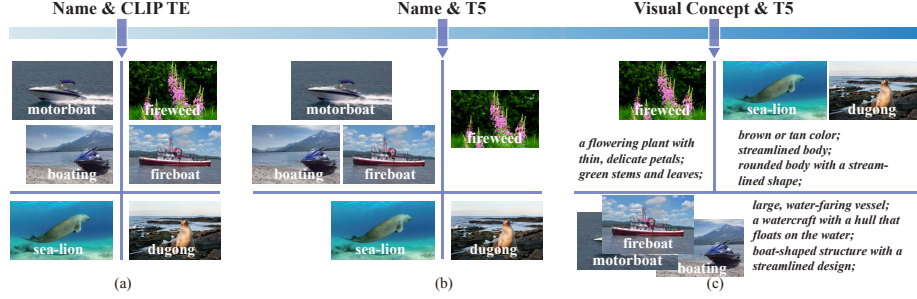
**Keywords:** Inter-category Relationships · Language Model · DETR

## 1 Introduction

Open-vocabulary object detection (OVOD) aims to identify and locate objects from a wide range of categories, including base and novel categories during inference, even though it is only trained on a limited set of base categories. Existing works [6, 9, 13, 29, 33, 35, 36, 40] in open-vocabulary object detection have been focusing on the development of sophisticated modules within detectors. These modules are tailored to effectively adapt the zero-shot and few-shot learning capabilities inherent in Vision-Language Models (VLMs) to the context of object detection.

---

\* Equal contribution. <sup>✉</sup> Corresponding author



**Fig. 1:** Illustration of the concept representation challenge. The clustering results are from (a) name embeddings by CLIP text encoder, (b) name embeddings by T5, and (c) visual description embeddings by T5, respectively. (a) CLIP text encoder struggles to distinguish between category names that are compositionally similar in letters, such as "fireboat" and "fireweed". (b) T5 fails to cluster categories that are visually comparable but compositionally different in name around the same cluster center, such as "sea-lion" and "dugong". (c) Marrying T5's textual semantic knowledge with visual insights achieves reasonable cluster results.

However, there are two challenges in most existing methods: (1) *Concept Representation*. Most existing methods represent concepts using name embeddings from CLIP text encoder. However, this approach of concept representation has a limitation in capturing the textual and visual semantic similarities between categories, which could aid in discriminating visually confusable categories and exploring potential novel objects; (2) *Overfit to base categories*. Although VLMs can perform well on novel categories, only base detection data is used in open vocabulary detectors' optimization, resulting in detectors' overfitting to base categories. As a result, novel objects are easily regarded as background or base categories.

Firstly, the issue of concept representation. Category names within CLIP's textual space are deficient in both textual depth and visual information.

(1) The VLM's text encoder lacks textual semantic knowledge compared with language model. As depicted in Figure 1(a), relying solely on name representations from CLIP concentrates on the similarity of letter composition, neglecting the hierarchical and common-sense understanding behind language. This method is disadvantageous for categorizing clustering as it fails to consider the conceptual relationships between categories. (2) Existing concept representations based on abstract category names or definitions fail to account for visual characteristics. Figure 1(b) demonstrates this problem, where sea lions and dugongs, despite their visual similarity, are allocated to separate clusters. Representing concept only with category name overlooks the rich visual context that language provides, which can facilitate the discovery of potential novel objects.

Secondly, the issue of overfitting to base categories. To leverage the open vocabulary capabilities of VLMs, we employ a frozen CLIP image encoder as the backbone and utilize category embeddings from the CLIP text encoder as

classification weights. We regard that detector training should serve two main functions: firstly, to differentiate foreground from background; and secondly, to maintain the open vocabulary classification capability of CLIP. However, training solely on base category annotations, without incorporating additional strategies, often results in overfitting: novel objects are commonly misclassified as either background or base categories. This problem has been further elucidated in prior research [29, 32].

We pinpoint the exploration of inter-category relationships as pivotal in tackling the aforementioned challenges. By cultivating a nuanced understanding of these relationships, we can develop a concept representation method that integrates both textual and visual semantics. This approach can also identify visually similar categories, guiding the model to focus more on learning generalized foreground features and preventing overfitting to base categories. Consequently, in this paper, we introduce **LaMI-DETR** (*Frozen CLIP-based DETR with **Language Model Instruction***), a simple but effective DETR-based detector that leverages language model insights to extract inter-category relationships, aiming to solve the aforementioned challenges.

To tackle the concept representation, we first adopt the Instructor Embedding [31], a T5 language model, to re-evaluate category similarities. As we find that language models exhibit a more refined semantic space compared to the CLIP text encoder. As shown in Figure 1(b), "fireweed" and "fireboat" are categorized into separate clusters, mirroring human recognition more closely. Next, we introduce the use of GPT-3.5 [2] to generate visual descriptions for each category. This includes detailing aspects such as shape, color, and size, effectively converting these categories into visual concepts. Figure 1(c) shows that, with similar visual descriptions, sea lions and dugongs are now grouped into the same cluster. To mitigate the overfitting issue, we cluster visual concepts into groups based on visual description embeddings from T5. This clustering result enables the identification and sampling of negative classes that are visually different from ground truth categories in each iteration. This relaxes the optimization of classification and focuses the model on deriving more generalized foreground features rather than overfitting to base categories. Consequently, this approach enhances the model's generalizability by reducing overtraining on base categories while preserving CLIP image backbone's ability to categorize.

In summary, we introduce a novel approach, LaMI, to enhance base-to-novel generalization in OVOD. LaMI harnesses large language models to extract inter-category relationships, utilizing this information to sample easy negative categories and avoid overfitting to base categories, while also refining concept representations to enable effective classification between visually similar categories. We propose a simple but effective end-to-end LaMI-DETR framework, enabling the effective transfer of open vocabulary knowledge from pretrained VLMs to detectors. We demonstrate the superiority of our LaMI-DETR framework through rigorous testing on large vocabulary OVOD benchmark, including +7.8 AP<sub>r</sub> on OV-LVIS and +2.9 AP<sub>r</sub> on VG-dedup(fair comparison with OWL [20, 22]). Code is available at <https://github.com/eternaldolphin/LaMI-DETR>.

## 2 Related Work

**Open-vocabulary object detection (OVOD)** leverages the image and language alignment knowledge stored in image-level dataset, *e.g.*, Conceptual Captions [28], or large pre-trained VLMs, *e.g.*, CLIP [25], to incorporate the open-vocabulary information into object detectors. One group of OVOD utilizes large-scale image-text pairs to expand detection vocabulary [7, 19, 26, 41, 44–46]. However, based on VLMs’ proven strong zero-shot recognition abilities, most open-vocabulary object detectors leverage VLM-derived knowledge to handle open vocabularies. The methods for object detectors to obtain open vocabulary knowledge from VLM can be divided into three categories: pseudo labels [26, 40, 45], distillation [6, 9, 33, 35] or parameter transfer [15, 36]. Despite its utility, performances of these methods are arguably restricted by the teacher VLM, which is shown to be largely unaware of inter-category visual relationship. Our method is orthogonal to all the aforementioned approaches in the sense that it not only explicitly models region-word correspondences, but also leverages visual correspondences across categories to help localize novel categories, which greatly improves the performance, especially in the DETR-based architecture [3, 11, 42, 43].

**Zero-shot object detection (ZSD)** addresses the challenge of detecting novel, unseen classes by leveraging language features for generalization. Traditional approaches utilize word embeddings, such as GloVe [23], as classifier weights to project region features into a pre-computed text embedding space [1, 5]. This enables ZSD models to recognize unseen objects by their names during inference. However, the primary limitation of ZSD lies in its training on a constrained set of seen classes, failing to adequately align the vision and language feature spaces. Some methods attempt to mitigate this issue by generating feature representations of novel classes using Generative Adversarial Networks [8, 30] or through data augmentation strategies for synthesizing unseen classes [48]. Despite these efforts, ZSD still faces significant performance gaps compared to supervised detection methods, highlighting the difficulty in extending detection capabilities to entirely unseen objects without access to relevant resources.

**Large Language Model (LLM)** Language data has increasingly played a pivotal role in open-vocabulary research, with recent Large Language Models (LLMs) showcasing vast knowledge applicable across various Natural Language Processing tasks. Works such as [21, 24, 37] have leveraged language insights from LLMs to generate descriptive labels for visual categories, thus enriching VLMs without necessitating further training or labeling. Nonetheless, there are gaps in current methodologies: firstly, the potential of discriminative LLMs for enhancing VLMs is frequently overlooked; secondly, the inter-category relationships remain underexplored. We propose a novel, straightforward clustering approach that employs GPT and Instructor Embeddings to investigate visual similarities among concepts, addressing these oversights.

### 3 Method

In this section, we begin with an introduction to open-vocabulary object detection (OVOD) in Section 3.1. Following this, we describe our proposed architecture of LaMI-DETR, a straightforward and efficient OVOD baseline, detailed in Section 3.2. Finally, we provide a detailed explanation of Language Model Instruction (LaMI) in Section 3.3.

#### 3.1 Preliminaries

Given an image  $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$  as input to an open-vocabulary object detector, two primary outputs are typically generated: (1) Classification, wherein a class label,  $c_j \in \mathcal{C}_{\text{test}}$ , is assigned to the  $j^{\text{th}}$  predicted object in the image, with  $\mathcal{C}_{\text{test}}$  representing the set of categories targeted during inference. (2) Localization, which involves determining the bounding box coordinates,  $\mathbf{b}_j \in \mathbb{R}^4$ , that identify the location of the  $j^{\text{th}}$  predicted object. Following the framework established by OVR-CNN [41], there is a detection dataset,  $\mathcal{D}_{\text{det}}$ , comprising bounding box coordinates, class labels, and corresponding images, and addressing a category vocabulary,  $\mathcal{C}_{\text{det}}$ .

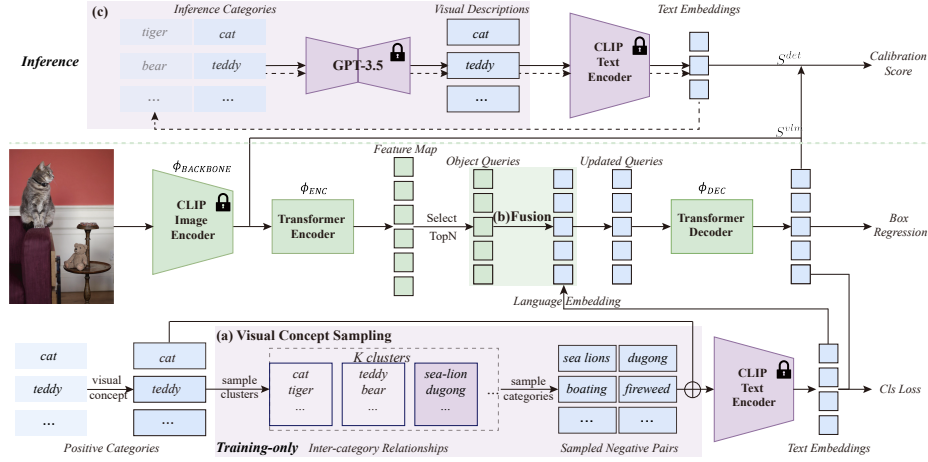
In line with the conventions of OVOD, we denote the category spaces of  $\mathcal{C}_{\text{test}}$  and  $\mathcal{C}_{\text{det}}$  as  $\mathcal{C}$  and  $\mathcal{C}_B$  respectively. Typically,  $\mathcal{C}_B \subset \mathcal{C}$ . The categories within  $\mathcal{C}_B$  are known as base categories, whereas those exclusively appearing in  $\mathcal{C}_{\text{test}}$  are identified as novel categories. The set of novel categories is expressed as  $\mathcal{C}_N = \mathcal{C} \setminus \mathcal{C}_B \neq \emptyset$ . For each category  $c \in \mathcal{C}$ , we utilize CLIP to encode its text embedding  $t_c \in \mathbb{R}^d$ , and  $\mathcal{T}_{\text{CLS}} = \{t_c\}_{c=1}^C$  ( $C$  is the size of the category vocabulary).

#### 3.2 Architecture of LaMI-DETR

The overall framework of LaMI-DETR is illustrated in Figure 2. Given an image input, we obtain the spatial feature map using the ConvNext backbone from the pre-trained CLIP image encoder ( $\Phi_{\text{BACKBONE}}$ ), which remains frozen during training. Then the feature map is sequentially subjected to a series of operations: a transformer encoder ( $\Phi_{\text{ENC}}$ ) to refine the feature map; a transformer decoder ( $\Phi_{\text{DEC}}$ ), producing a set of query features  $\{f_j\}_{j=1}^N$ ; The query features are then processed by a bounding box module ( $\Phi_{\text{BBOX}}$ ) to infer the positions of objects, denoted as  $\{\mathbf{b}_j\}_{j=1}^N$ . We follow the inference pipeline of F-VLM [15] and use VLM score  $S^{vlm}$  to calibrate detection score  $S^{det}$ .

$$S_j^{vlm} = \mathcal{T}_{\text{CLS}} \cdot \Phi_{\text{pooling}}(b_j) \quad (1)$$

$$S_c^{cal} = \begin{cases} S_c^{vlm^\alpha} \cdot S_c^{det^{(1-\alpha)}} & \text{if } c \in \mathcal{C}_B \\ S_c^{vlm^\beta} \cdot S_c^{det^{(1-\beta)}} & \text{if } c \in \mathcal{C}_N \end{cases} \quad (2)$$



**Fig. 2:** An overview of LaMI-DETR Framework. LaMI-DETR adapts the DETR model by incorporating a frozen CLIP image encoder as the backbone and replacing the final classification layer with CLIP text embeddings. (a) Visual Concept Sampling, applied only during the training phase, leverages pre-extracted inter-category relationships to sample easy negative categories that are visually distinct from ground truth classes. This encourages the detector to derive more generalized foreground features rather than overfitting to base categories. (b) Language embeddings selected are integrated into the object queries for enhanced classification accuracy. (c) During inference, confusing categories are identified to improve VLM score.

**Comparison with other Open-Vocabulary DETR.** CORA [36] and EdaDet [29] also propose to use a frozen CLIP image encoder in DETR for extracting image features. However, LaMI-DETR significantly differs from these two approaches in the following aspects.

Firstly, regarding the number of backbones used, both LaMI-DETR and CORA employ a single backbone. In contrast, EdaDet utilizes two backbones: a learnable backbone and a frozen CLIP image encoder.

Secondly, both CORA and EdaDet adopt an architecture that decouples classification and regression tasks. While this method addresses the issue of failing to recall novel classes, it necessitates extra post-processing steps, such as NMS, disrupting DETR’s original end-to-end structure.

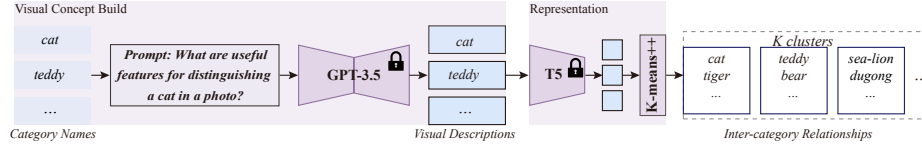
Furthermore, both CORA and EdaDet require RoI-Align operations during training. In CORA, the DETR only predicts objectness, necessitating RoI-Align on the CLIP feature map during anchor pre-matching to determine the specific categories of proposals. EdaDet minimizes the cross-entropy loss based on each proposal’s classification scores, obtained through a pooling operation. Consequently, CORA and EdaDet require multiple pooling operations during inference. In contrast, LaMI-DETR simplifies this process, needing only a single pooling operation at the inference stage.

### 3.3 Language Model Instruction

Unlike previous methods that only rely on the vision-language alignment of VLMs, we aim to improve open-vocabulary detectors by enhancing concept representation and investigating inter-category relationships. To achieve this, we first explain the process of constructing visual concepts and delineating their relationships. In *Language Embedding Fusion* and *Confusing Category* sections, we describe methods for more accurately representing concepts during the training and inference processes. The *Visual Concept Sampling* section addresses how to mitigate overfitting issue through the use of inter-category relationships. Finally, we detail the distinctions with other research effort.

**Inter-category Relationships Extraction.** Based on the problem identified in Figure 1, we employ visual descriptions to establish visual concepts, refining concept representation. Furthermore, we utilize T5, which possesses extensive textual semantic knowledge, to measure similarity relationships among visual concepts, thereby extracting inter-category relationships.

As illustrated in Figure 3, given a category name  $c \in \mathcal{C}$ , we extract its fine-grained visual feature descriptors  $d$  using the method described in 21. We define  $\mathcal{D}$  as the visual description space for categories in  $\mathcal{C}$ . These visual descriptions  $d \in \mathcal{D}$  are then sent to the T5 model to obtain the visual description embeddings  $e \in \mathcal{E}$ . Consequently, we construct an open set of visual concepts  $\mathcal{D}$  and their corresponding embeddings  $\mathcal{E}$ . To identify visually similar concepts, we propose clustering the visual description embeddings  $\mathcal{E}$  into  $K$  cluster centroids. Concepts grouped under the same cluster centroid are deemed to possess similar visual characteristics. The extracted inter-category relationships are then applied in the visual concept sampling as shown in Figure 2(a).



**Fig. 3:** Illustration of Inter-category Relationships Extraction. Visual descriptions generated by GPT-3.5 are processed by T5 to cluster categories with visual similarities.

**Language Embedding Fusion.** As shown in Figure 2(b), after transformer encoder, each pixel on the feature map  $\{f_i\}_{i=1}^M$  is interpreted as an object query, with each directly predicting a bounding box. To select the top  $N$  scoring bounding boxes as region proposals, the process can be encapsulated as follows:

$$\{q_j\}_{j=1}^N = \text{Top}_N(\{\mathcal{T}_{\text{CLS}} \cdot f_i\}_{i=1}^M). \quad (3)$$

In LaMI-DETR, we fuse each query  $\{q_j\}_{j=1}^N$  with its closest text embedding, resulting in:

$$\{q_j\}_{j=1}^N = \{q_j \oplus t_j\}_{j=1}^N, \quad (4)$$

where  $\oplus$  denotes element-wise addition.

On one hand, the visual descriptions are sent to the T5 model to cluster visually similar categories, as previously described. On the other hand, the visual descriptions  $d_j \in \mathcal{D}$  are forwarded to the text encoder of the CLIP model to update the classification weights, denoted as  $\mathcal{T}_{\text{CLS}} = \{t'_c\}_{c=1}^C$ , where  $t'_c$  represents the text embedding of  $d$  in the CLIP text encoder space. Consequently, the text embeddings used in the language embedding fusion process are updated accordingly:

$$\{q_j\}_{j=1}^N = \{q_j \oplus t'_j\}_{j=1}^N \quad (5)$$

**Confusing Category.** Due to similar visual concepts often sharing common features, nearly identical visual descriptors can be generated for these categories. This similarity poses challenges in distinguishing similar visual concepts during the inference process.

To distinguish easily confusable categories during the inference process, we initially identify the most similar category  $c^{\text{conf}} \in \mathcal{C}$  for each class  $c \in \mathcal{C}$  within the CLIP text encoder semantic space, based on  $\mathcal{T}_{\text{CLS}}$ . We then modify the prompt for generating visual descriptions  $d' \in \mathcal{D}'$  for category  $c$  to emphasize the features that differentiate  $c$  from  $c^{\text{conf}}$ . Let  $t''$  be the text embedding of  $d'$  in the CLIP text encoder space. As shown in Figure 2(c), We update the inference pipeline as follows:

$$\mathcal{T}'_{\text{CLS}} = \{t''_c\}_{c=1}^C, \quad (6)$$

$$S_j^{\text{vlm}} = \mathcal{T}'_{\text{CLS}} \cdot \Phi_{\text{pooling}}(b_j). \quad (7)$$

**Visual Concept Sampling.** To address the challenges posed by incomplete annotations in open-vocabulary detection datasets, we employ Federated Loss [47], originally introduced for long-tail datasets [10]. This approach involves randomly selecting a set of categories to calculate detection losses for each mini-batch, effectively minimizing issues related to missing annotations in certain classes. Given category occurrence frequency  $p = [p_1, p_2, \dots, p_C]$ , where  $p_c$  denotes the occurrence frequency in training data of the  $c^{\text{th}}$  visual concept and  $C$  represents the total number of categories. We randomly draw  $C_{\text{fed}}$  samples based on the probability distribution  $p$ . The likelihood of selecting the  $c^{\text{th}}$  sample  $x_c$  is proportional to its corresponding weight  $p_c$ . This method facilitates the transfer of visual similarity knowledge, extracted by the language model, to the detector, thereby reducing the issue of overfitting:

$$P(X = c) = p_c, \quad \text{for } c = 1, 2, \dots, C \quad (8)$$

Incorporating federated loss, the classification weight is reformulated as  $\mathcal{T}_{\text{CLS}} = \{t''_c\}_{c=1}^{C_{\text{FED}}}$ , where  $C_{\text{FED}}$  denotes the categories engaged in the loss calculation of each iteration, and  $C_{\text{fed}}$  is the count of  $C_{\text{FED}}$ .

We utilize a frozen CLIP with strong open vocabulary capabilities as LaMI-DETR’s backbone. However, due to the limited categories in detection datasets, overfitting to base classes is inevitable after training. To mitigate overtraining on base categories, we aim to sample straightforward negative categories based on the results of visual concepts clustering. In LaMI-DETR, let the clusters containing the ground truth categories be denoted by  $\mathcal{K}_G$  in a given iteration. We denote all the categories within  $\mathcal{K}_G$  as  $\mathcal{C}_g$ . Specifically, we aim to exclude  $\mathcal{C}_g$  from being sampled in the current iteration. To achieve this, we set the frequency of occurrence for categories within  $\mathcal{C}_g$  to zero. This approach enables the transfer of visual similarity knowledge, extracted by the language model, to the detector, mitigating overfitting issue:

$$p_c^{\text{cal}} = \begin{cases} 0 & \text{if } c \in \mathcal{C}_g \\ p_c & \text{if } c \notin \mathcal{C}_g \end{cases} \quad (9)$$

where  $p_c^{\text{cal}}$  indicates the frequency of occurrence of category  $c$  after language model calibration, ensuring visually similar categories are not sampled during this iteration. This process is shown in Figure 2(a).

**Comparison with concept enrichment.** The visual concept description is different from the concept enrichment employed in DetCLIP [39]. The visual descriptions used in LaMI place more emphasis on the visual attributes inherent to the object itself. In DetCLIP, category label is supplemented with definitions, which may include concepts not present in the pictures to rigorously characterize a class.

## 4 Experiments

Section 4.1 introduces the standard dataset and benchmarks commonly utilized in the field, as detailed in [9]. Section 4.2 outlines the implementation and training details of our LaMI-DETR, which leverages knowledge of visual characteristics from language models. We present a comparison of our models with existing works in Section 4.3, showcasing state-of-the-art performance. Additionally, Section 4.3 includes results on cross-dataset transfer to demonstrate the generalizability of our approach. Finally, Section 4.4 conducts ablation studies to examine the impact of our design decisions.

### 4.1 Datasets

**LVIS.** Our experiments are conducted on the LVIS dataset, which includes annotations for 1,203 object categories. These categories are divided into three

groups—rare, common, and frequent—based on the number of training images containing a given class. Following the approach of previous studies, we categorize them into 866 base classes, encompassing frequent and common categories, and 337 novel classes, consisting of rare categories. To create an open-vocabulary scenario, we exclude annotations for novel classes from the training images. In line with standard practice, we report the mean average precision (mAP) for predicted boxes specifically for the rare classes, denoted as  $AP_r$ . Additionally, we present the box AP averaged across all classes to reflect overall performance, denoted as mAP.

**Object365 and VisualGenome.** For a fair comparison with OWL-ViT [20, 22], we adopt the same training settings, utilizing data from Object365 and VisualGenome. To conserve training time, we employ only a 1/3 random sample of Object365 in our study. With respect to VisualGenome, we meticulously replicate OWL-ViT’s preprocessing steps by eliminating all detection annotations that correspond to the names of LVIS’s rare categories. The resulting curated dataset is referred to as VG dedup.

## 4.2 Implementation Details

Training is conducted on 8 40G A100 GPUs with a total batch size of 32. For the OV-LVIS setting, we train the model for 12 epochs. In the VG-dedup benchmark, to ensure a fair comparison with OWL-ViT, we initially pretrain LaMI-DETR on a randomly sampled 1/3 of the Object365 dataset for 12 epochs. Subsequently, LaMI-DETR is finetuned on the VG dedup dataset for an additional 12 epochs.

The detector utilizes ConVNext-Large [18] from OpenCLIP [12] as its backbone, which remains frozen throughout the training process. LaMI-DETR, building upon DINO, employs 900 queries as specified in detrex [27]. We adhere closely to the original training configurations detailed in detrex, with the exception of employing an exponential moving average (EMA) strategy to enhance training stability. To balance the distribution of training samples, we apply repeat factor sampling [10] using the default hyperparameters. For federated loss, the numbers of categories  $C_{\text{fed}}$  are set to 100 and 700 for OV-LVIS and VG dedup datasets, respectively.

To explore a broader range of visual concepts for more effective clustering, we compile a comprehensive category collection from LVIS, Object365, VisualGenome, Open Images, and ImageNet-21K. Redundant concepts are filtered out using WordNet hypernyms, resulting in a visual concept dictionary comprising 26,410 unique concepts. During the visual concept grouping phase, this dictionary is clustered into  $K$  centers, with  $K$  being 128 for OV-LVIS and 256 for VG dedup, respectively.

## 4.3 Open-Vocabulary Detection Results

**Table 1: LVIS open-vocabulary detection (box AP).** LaMI-DETR outperforms the best existing approach by +7.8 box AP<sub>r</sub> in the standard benchmark. All methods use the same instance-level supervision from LVIS base categories for detection training. †: reports mask AP. \*: uses the image-level data in pretraining. We calculate the backbone’s parameters based on models released by CLIP except RN50, which may vary slightly from their actual sizes.

Method	Pretrained Model	Detector Backbone	Backbone Size	Image-level Dataset	AP <sub>r</sub>	AP
VL-PLM [44]	ViT-B/32	R-50	26M	IN-L	17.2 <sup>†</sup>	27.0 <sup>†</sup>
OV-DETR [40]	ViT-B/32	R-50	26M	✗	17.4 <sup>†</sup>	26.6 <sup>†</sup>
DetPro-Cascade [6]	ViT-B/32	R-50	26M	✗	21.7	30.5
Rasheed [26]	ViT-B/32	R-50	26M	IN-L	21.1 <sup>†</sup>	25.9 <sup>†</sup>
PromptDet [7]	ViT-B/32	R-50	26M	LAION-novel	21.4 <sup>†</sup>	25.3 <sup>†</sup>
OADP [33]	ViT-B/32	R-50	26M	✗	21.9	28.7
RegionCLIP [45]	R-50x4	R-50x4	87M	CC3M	22.0 <sup>†</sup>	32.3 <sup>†</sup>
CORA [36]	R-50x4	R-50x4	87M	✗	22.2	-
BARON [35]	ViT-B/32	R-50	26M	CC3M	23.2	29.5
CondHead [34]	R-50x4	R-50x4	87M	CC3M	25.1	33.7
Detic-CN2 [46]	ViT-B/32	R-50	26M	IN-L	24.6 <sup>†</sup>	32.4 <sup>†</sup>
ViLD-Ens [9]	ViT-B/32	R-50	26M	✗	16.7	27.8
F-VLM [15]	R-50x64	R-50x64	420M	✗	32.8 <sup>†</sup>	34.9 <sup>†</sup>
OWL-ViT [22]	ViT-L/14	ViT-L/14	306M	✗	25.6	34.7
RO-ViT [14]	ViT-B/16	ViT-B/16	86M	ALIGN*	28.4	31.9
RO-ViT [14]	ViT-L/16	ViT-L/16	303M	ALIGN*	33.6	36.2
CFM-ViT [13]	ViT-B/16	ViT-B/16	86M	ALIGN*	29.6	33.8
CFM-ViT [13]	ViT-L/16	ViT-L/16	303M	ALIGN*	35.6	38.5
<b>ours</b>	ConVNext-L	ConVNext-L	196M	✗	<b>43.4</b>	<b>41.3</b>

**OV-LVIS.** We compare our LaMI-DETR framework with the other state-of-the-art OVOD methods in Table 1. We report overall box AP performance and box AP for "rare" classes only. The latter metric is the key measure of OVOD performance. Our method obtain the best performance on both AP<sub>r</sub> and overall mAP compared to existing approaches for open-vocabulary object detection, while utilizing a more challenging strictly open-vocabulary training paradigm without additional data. LaMI-DETR, with a backbone of only 196M parameters significantly less than CFM-ViT’s 303M achieves superior performance. Moreover, LaMI-DETR does not utilize additional image-level datasets. The results demonstrate that LaMI-DETR has lower computational requirements and higher accuracy.

**Zero-shot LVIS.** We evaluate the model’s ability to recognize diverse and rare objects on LVIS in a zero-shot setting. We replace VG-dedup with LVIS vocabulary embeddings for zero-shot detection without finetuning. We assume all categories are novel and set  $\alpha, \beta = (0.0, 0.25)$  in Eq 2. We use OWL as the baseline

**Table 2: LVIS zero-shot detection (box AP).** §: The models only report fixed AP [4] on *LVIS-val*. The models depicted in this figure utilize multiple detection datasets, excluding LVIS; therefore, we refer to this configuration as the zero-shot setting.

Method	Detector Backbone	Datasets	<b>AP<sub>r</sub></b> AP	
GLIP-L [16]	Swin-L	O365,GoldG,Cap4M	17.1	26.9
GroundingDINO [17]	Swin-L	O365,GoldG,OI,Cap4M,COCO,RefC	22.0	32.3
DetCLIP <sup>§</sup> [39]	Swin-L	O365,GoldG,YFCC1M	27.6	31.2
DetCLIPv2 <sup>§</sup> [38]	Swin-L	O365,GoldG,CC15M	33.3	36.6
OWL-ViT [22]	ViT-L/14	O365,VG-dedup	31.2	34.6
OWL-ST [20]	ViT-L/14	O365,VG-dedup	34.9	33.5
<b>ours</b>	ConVNext-L	O365,VG-dedup	<b>37.8</b>	35.4

for our models. The results are shown in Table 2. LaMI-DETR outperforms OWLs under the same settings.

**Cross-dataset Transfer.** To evaluate the generalizability of our method in a cross-dataset transfer detection setting, we conduct experiments on the COCO and Objects365-v1 validation split. Specifically, we directly applies the detector trained on the LVIS base categories, while replacing the LVIS class embeddings with those of COCO/Objects365 for transfer detection without further finetuning. All categories were treated as novel. Our best-performing model achieved 42.8 AP on COCO and 21.9 AP on Object365, outperforming CoDet by +3.7 AP on COCO and CFM by +3.2 AP on Object365 according to Table 3.

**Table 3:** Cross-datasets transfer detection from OV-LVIS to COCO and Objects365. F-VLM adopts RN50 in CLIP as backbone, which is larger than standard RN50.

Method	Backbone	Parameters	COCO			Objects365		
			AP	AP <sub>50</sub>	AP <sub>75</sub>	AP	AP <sub>50</sub>	AP <sub>75</sub>
ViLD [9]	RN50	26M	36.6	55.6	39.8	11.8	18.2	12.6
DetPro [6]	RN50	26M	34.9	53.8	37.4	12.1	18.8	12.9
F-VLM [15]	RN50	38M	32.5	53.1	34.6	11.9	19.2	12.6
BARON [35]	RN50	26M	36.2	55.7	39.1	13.6	21.0	14.5
CoDet [19]	EVA02-L	304M	39.1	57.0	42.3	14.2	20.5	15.3
CFM [13]	ViT-L/16	303M	-	-	-	18.7	28.9	20.3
<b>ours</b>	ConvNext-L	196M	<b>42.8</b>	<b>57.6</b>	<b>46.9</b>	<b>21.9</b>	<b>30.0</b>	<b>23.5</b>

**Table 4:** Ablations for our model. Language Model Instruction consists of visual concepts sampling, embedding update and confusing categories distinguish. Below the horizontal line are the results with the class factor. See Table 7 for details.

#	Federated Loss(Eq.8)	Embedding Fusion(Eq.4)	Visual Concepts Sampling(Eq.9)	Embedding Update(Eq.5)	Confusing Category(Eq.6)	AP <sub>r</sub>
1	✓					32.2
2	✓	✓				33.0
3	✓	✓	✓			40.1
4	✓	✓	✓	✓		42.5
5	✓	✓	✓	✓	✓	43.4

#### 4.4 Ablation Study

To study the advantages of LaMI-DETR, we provide ablation studies on the OV-LVIS benchmark.

**LaMI-DETR.** Table 4 demonstrates the impact of incorporating language model guidance into our LaMI-DETR framework. The version without LaMI module achieves an AP<sub>r</sub> of 33.0. By integrating our proposed LaMI module, the model achieves an AP<sub>r</sub> of 43.4. The top two rows in Table 4 shows language embedding fusion in Eq.4 brings a 0.8 AP<sub>r</sub> gain. The 3<sup>rd</sup> to 5<sup>th</sup> row in Table 4 adds Visual Concept Sampling, embedding update and Confusing Category distinguishing to baseline gradually.

**Table 5:** Ablation study on the confusing category. Zero-shot proposal classification performance on LVIS minival datasets.

Model	mAcc <sub>r</sub>	mAcc <sub>c</sub>	mAcc <sub>f</sub>	mAcc
CLIP	43.8	44.1	37.8	41.0
visual desc. [21]	49.5	45.8	40.2	43.4
ours	52.7	46.1	41.4	44.4

**Confusing Category.** We demonstrate the effectiveness of Confusing Category in Table 5. Given the ground truth bounding boxes, we use different text embeddings to classify their region features. To evaluate the performance, we compute "Mean Accuracy" (accuracy for each category independently with equal weights). For the following strategies, we use RoI-Align to directly extract features from CLIP. The table validates that the CLIP text encoder can discriminate categories from confusing ones with our refined concept representation.

**Table 6:** Ablation study on the cluster designs. For fair comparison, all detectors use classification weights from CLIP text encoder name embeddings. †: Results with class factor. See Table 7 for details.

Model	Cluster Encoder	Cluster Text	AP <sub>r</sub>	AR <sub>r</sub>
baseline	-	-	33.0	40.3
baseline+VCS	CLIP Text Encoder	name	33.5	41.4
baseline+VCS	Instructor Embedding	name	34.1	39.5
baseline+VCS	Instructor Embedding	name+definition	31.5	37.3
baseline+VCS	Instructor Embedding	name+visual desc.	40.1 <sup>†</sup>	57.0 <sup>†</sup>

**The Cluster Design.** Visual Concept Sampling aims to sample negative categories with large visual differences from the positive class, enabling the detector to utilize inter-class relationships by penalizing categories with large visual differences, thus achieving generalization to visually close classes. We validate this claim through the enhancements detailed in Table 6.

Our experimental results in Table 6 demonstrate the effectiveness of our sampling negative classes method. The first row shows results of baseline. Rows 2<sup>nd</sup>-5<sup>th</sup> employ the Visual Concept Sampling module but vary the clustering method. Specifically, the second row clusters category name embeddings from a CLIP text encoder, corresponding to (a) in Figure 1. The third row clusters category name embeddings from the T5 space, corresponding to Figure 1(b). The fourth row aims to match DetCLIP’s concept enrichment by clustering definition embeddings in the T5 space. Finally, the last row presents our full method, clustering category visual description embeddings from the T5 space as Figure 1(c). This systematic ablation analyzes how different semantics and grouping strategies within the Visual Concept Sampling module affect downstream detection performance, validating the importance of visual similarity-based concept sampling for our task.

## 5 Conclusion

In this paper, we undertake the first effort to explore inter-category relationships for generalization in OVOD. We introduce LaMI-DETR, a framework that effectively utilize the visual concepts similarity to sample negative categories during training for learning generalizable object localization and retaining open vocabulary knowledge of VLMs. Additionally, the refined concepts enable effective object classification especially between confusing categoris. Experiments show that LaMI-DETR achieves state-of-the-art performance across various OVOD benchmarks. On the other hand, our method utilizes the CLIP ConvNext-L architecture as the visual backbone. Exploring alternative pre-trained VLMs such as those based on ViT is under-explored here. We leave this for further investigation.

## Acknowledgements

This research is supported in part by National Science and Technology Major Project (2022ZD0115502), National Natural Science Foundation of China (NO. 62122010, U23B2010), Zhejiang Provincial Natural Science Foundation of China (Grant No. LDT23F02022F02), and Beijing Natural Science Foundation (NO. L231011). We thank the authors of LW-DETR [3]: Qiang Chen and Xinyu Zhang, the author of OADP [33]: Yi Liu and the author of DetPro [6]: Yu Du for their helpful discussions.

## References

1. Bansal, A., Sikka, K., Sharma, G., Chellappa, R., Divakaran, A.: Zero-shot object detection. In: ECCV (2018)
2. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. NIPS **33**, 1877–1901 (2020)
3. Chen, Q., Su, X., Zhang, X., Wang, J., Chen, J., Shen, Y., Han, C., Chen, Z., Xu, W., Li, F., et al.: Lw-detr: A transformer replacement to yolo for real-time detection. arXiv preprint arXiv:2406.03459 (2024)
4. Dave, A., Dollár, P., Ramanan, D., Kirillov, A., Girshick, R.: Evaluating large-vocabulary object detectors: The devil is in the details. arXiv preprint arXiv:2102.01066 (2021)
5. Demirel, B., Cinbis, R.G., Ikizler-Cinbis, N.: Zero-shot object detection by hybrid region embedding. In: BMVC (2018)
6. Du, Y., Wei, F., Zhang, Z., Shi, M., Gao, Y., Li, G.C.: Learning to prompt for open-vocabulary object detection with vision-language model. CVPR (2022)
7. Feng, C., Zhong, Y., Jie, Z., Chu, X., Ren, H., Wei, X., Xie, W., Ma, L.: Promptdet: Towards open-vocabulary detection using uncurated images. In: ECCV. pp. 701–717. Springer (2022)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014)
9. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. In: ICLR (2021)
10. Gupta, A., Dollar, P., Girshick, R.: LVIS: A dataset for large vocabulary instance segmentation. In: CVPR (2019)
11. Hu, Z., Sun, Y., Wang, J., Yang, Y.: Dac-detr: Divide the attention layers and conquer. Advances in Neural Information Processing Systems **36** (2024)
12. Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., Hajishirzi, H., Farhadi, A., Schmidt, L.: Openclip (Jul 2021). <https://doi.org/10.5281/zenodo.5143773> <https://doi.org/10.5281/zenodo.5143773> if you use this software, please cite it as below.
13. Kim, D., Angelova, A., Kuo, W.: Contrastive feature masking open-vocabulary vision transformer (2023)
14. Kim, D., Angelova, A., Kuo, W.: Region-aware pretraining for open-vocabulary object detection with vision transformers. In: CVPR. pp. 11144–11154 (2023)
15. Kuo, W., Cui, Y., Gu, X., Piergiovanni, A., Angelova, A.: Open-vocabulary object detection upon frozen vision and language models. In: ICLR (2023), <https://openreview.net/forum?id=MIMwy4kh91f>

16. Li\*, L.H., Zhang\*, P., Zhang\*, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., Chang, K.W., Gao, J.: Grounded language-image pre-training. In: CVPR (2022)
17. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023)
18. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. CVPR (2022)
19. Ma, C., Jiang, Y., Wen, X., Yuan, Z., Qi, X.: Codet: Co-occurrence guided region-word alignment for open-vocabulary object detection. In: NIPS (2023)
20. Matthias Minderer, Alexey Gritsenko, N.H.: Scaling open-vocabulary object detection. NeurIPS (2023)
21. Menon, S., Vondrick, C.: Visual classification via description from large language models. ICLR (2023)
22. Minderer, M., Gritsenko, A., Stone, A., Neumann, M., Weissenborn, D., Dosovitskiy, A., Mahendran, A., Arnab, A., Dehghani, M., Shen, Z., et al.: Simple open-vocabulary object detection. In: ECCV. pp. 728–755. Springer (2022)
23. Pennington, J., Socher, R., Manning, C.D.: GloVe: Global vectors for word representation. In: EMNLP (2014)
24. Pratt, S., Covert, I., Liu, R., Farhadi, A.: What does a platypus look like? generating customized prompts for zero-shot image classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15691–15701 (2023)
25. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763. PMLR (2021)
26. Rasheed, H., Maaz, M., Khattak, M.U., Khan, S., Khan, F.S.: Bridging the gap between object and image-level representations for open-vocabulary detection. In: NIPS (2022)
27. Ren, T., Liu, S., Li, F., Zhang, H., Zeng, A., Yang, J., Liao, X., Jia, D., Li, H., Cao, H., Wang, J., Zeng, Z., Qi, X., Yuan, Y., Yang, J., Zhang, L.: detrex: Benchmarking detection transformers (2023)
28. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2556–2565 (2018)
29. Shi, C., Yang, S.: Edadet: Open-vocabulary object detection using early dense alignment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2023)
30. Shizhen, Z., Changxin, G., Yuanjie, S., Lerenhan, L., Changqian, Y., Zhong, J., Nong, S.: Gtnet: Generative transfer network for zero-shot object detection. In: AAAI (2020)
31. Su, H., Shi, W., Kasai, J., Wang, Y., Hu, Y., Ostendorf, M., Yih, W.t., Smith, N.A., Zettlemoyer, L., Yu, T.: One embedder, any task: Instruction-finetuned text embeddings (2022), <https://arxiv.org/abs/2212.09741>
32. Wang, J., Zhang, H., Hong, H., Jin, X., He, Y., Xue, H., Zhao, Z.: Open-vocabulary object detection with an open corpus. In: ICCV. pp. 6759–6769 (2023)
33. Wang, L., Liu, Y., Du, P., Ding, Z., Liao, Y., Qi, Q., Chen, B., Liu, S.: Object-aware distillation pyramid for open-vocabulary object detection. CVPR (2023)
34. Wang, T.: Learning to detect and segment for open vocabulary object detection. In: CVPR. pp. 7051–7060 (2023)

35. Wu, S., Zhang, W., Jin, S., Liu, W., Loy, C.C.: Aligning bag of regions for open-vocabulary object detection. In: CVPR (2023)
36. Wu, X., Zhu, F., Zhao, R., Li, H.: Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. ArXiv **abs/2303.13076** (2023)
37. Yang, Y., Panagopoulou, A., Zhou, S., Jin, D., Callison-Burch, C., Yatskar, M.: 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. pp. 19187–19197 (2023)
38. Yao, L., Han, J., Liang, X., Xu, D., Zhang, W., Li, Z., Xu, H.: 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. pp. 23497–23506 (2023)
39. Yao, L., Han, J., Wen, Y., Liang, X., Xu, D., Zhang, W., Li, Z., Xu, C., Xu, H.: Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. NIPS **35**, 9125–9138 (2022)
40. Zang, Y., Li, W., Zhou, K., Huang, C., Loy, C.C.: Open-vocabulary detr with conditional matching (2022)
41. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: CVPR. pp. 14393–14402 (2021)
42. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection (2022)
43. Zhao, C., Sun, Y., Wang, W., Chen, Q., Ding, E., Yang, Y., Wang, J.: Ms-detr: Efficient detr training with mixed supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17027–17036 (2024)
44. Zhao, S., Zhang, Z., Schuster, S., Zhao, L., Vijay Kumar, B., Stathopoulos, A., Chandraker, M., Metaxas, D.N.: Exploiting unlabeled data with vision and language models for object detection. In: ECCV. pp. 159–175. Springer (2022)
45. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. In: CVPR. pp. 16793–16803 (2022)
46. Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twenty-thousand classes using image-level supervision. In: ECCV. pp. 350–368. Springer (2022)
47. Zhou, X., Koltun, V., Krähenbühl, P.: Probabilistic two-stage detection. In: arXiv preprint arXiv:2103.07461 (2021)
48. Zhu, P., Wang, H., Saligrama, V.: Don’t Even Look Once: Synthesizing features for zero-shot detection. In: CVPR (2020)