

Comprehensive Multi-Modal Prototypes Are Simple and Effective Classifiers for Vast-Vocabulary Object Detection

Yitong Chen^{1,2*}, Wenhao Yao^{1*}, Lingchen Meng^{1*}, Sihong Wu¹, Zuxuan Wu^{1,2†}, Yu-Gang Jiang¹

¹Shanghai Key Lab of Intell. Info. Processing, School of CS, Fudan University

²Shanghai Innovation Institute

{chenyt24, whyao23, shwu22}@m.fudan.edu.cn, {lcmeng20, zxwu, ygj}@fudan.edu.cn

Abstract

Enabling models to recognize vast open-world categories has been a longstanding pursuit in object detection. By leveraging the generalization capabilities of vision-language models, current open-world detectors can recognize a broader range of vocabularies, despite being trained on limited categories. However, when the scale of the category vocabularies during training expands to a real-world level, previous classifiers aligned with coarse class names significantly reduce the recognition performance of these detectors. In this paper, we introduce **Prova**, a multi-modal prototype classifier for vast-vocabulary object detection. Prova extracts comprehensive multi-modal prototypes as initialization of alignment classifiers to tackle the vast-vocabulary object recognition failure problem. On V3Det, this simple method greatly enhances the performance among one-stage, two-stage, and DETR-based detectors with only additional projection layers in both supervised and open-vocabulary settings. In particular, Prova improves Faster R-CNN, FCOS, and DINO by 3.3, 6.2, and 2.9 AP respectively in the supervised setting of V3Det. For the open-vocabulary setting, Prova achieves a new state-of-the-art performance with 32.8 base AP and 11.0 novel AP, which is of 2.6 and 4.3 gain over the previous methods.

Code — <https://github.com/Row11n/Prova>

Introduction

The ability of humans to recognize a vast number of object categories in real-world scenarios has always been the goal of open-world object detection. Although fully supervised closed-set object detection methods have made significant progresses (Ren et al. 2015; Tian et al. 2022; Wang et al. 2021; Zhu et al. 2021; Zhang et al. 2023a), they still struggle to handle vast categories in the wild due to the limited number of object categories in commonly used detection datasets, e.g., COCO (Lin et al. 2014) with 80 categories and Object365 (Shao et al. 2019) with 365 categories. To address this issue, remarkable advancements have been achieved in an open-vocabulary manner which have acceptable zero-shot ability to unseen categories during training and also comparable closed-set capability (Zhou et al. 2022; Zhong et al.

*These authors contributed equally.

†Corresponding author.

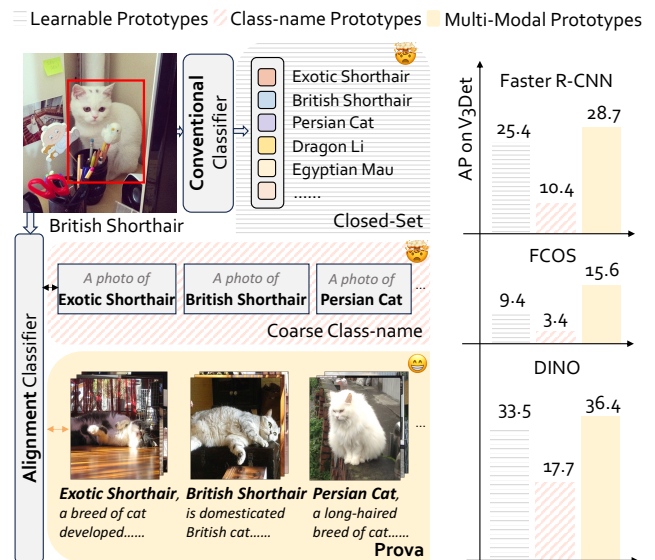


Figure 1: Left: When the category vocabularies scale up, previous conventional closed-set classifiers and class name alignment classifiers struggle to distinguish complex classes such as *Exotic Shorthair* and *British Shorthair*. Our key idea is to extract more detailed multi-modal prototypes to replace coarse class names with templates in the alignment classifier. Right: AP on V3Det validation set in a supervised manner. The class name alignment classifier performs even worse than the conventional classifier, and multi-modal prototype classifier, *i.e.* Prova, achieves the best results.

2022). Most of the recent open-vocabulary detectors leverage the generalization capabilities of Vision-Language Models (VLMs) like CLIP (Radford et al. 2021), by converting the closed-set prediction head into vision-language alignment between objects and class names.

However, as shown in Figure 1, when the number of categories scales up in the detection dataset, e.g. V3Det (Wang et al. 2023) with 13204 categories, the class name alignment classifier, which is well done in open-vocabulary fields, fails and performs even worse than the conventional closed-set classifier. Previous vision-language alignment methods assume a mapping relationship from detected regions to class

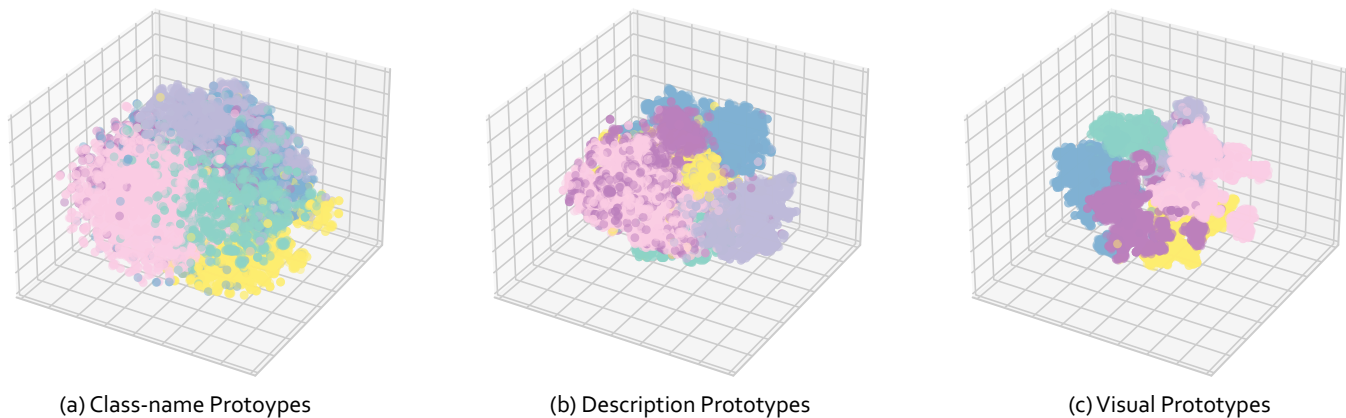


Figure 2: Feature space visualization of (a) Class names with fixed templates; (b) Detailed class descriptions; (c) Reference images of each class. We use CLIP-ViT-Large (Radford et al. 2021) as the text encoder to extract class name prototypes, LongCLIP-Large (Zhang et al. 2024) to extract description prototypes, CLIP-ViT-Large visual encoder to extract visual prototypes and then utilize t-SNE (van der Maaten and Hinton 2008) after k -means. Compared to class name prototypes, description and visual prototypes have a cleaner feature space, which is preferable for alignment classifier.

vocabularies (Gu et al. 2022; Zhong et al. 2022), which holds when the vocabulary size is small, but becomes problematic as the vocabulary size expands. Specifically, coarse class names with fixed templates as the textual guidance encounter issues like semantic confusion (e.g., *Exotic Shorthair* and *British Shorthair* shown in Figure 1) and ambiguity (e.g., *Crane* in V3Det categories, which could be a machine or an animal belonging to Gruiformes), leading to difficulties in maintaining the mapping relation.

Furthermore, it is challenging even for a human to distinguish whether the cat detected in Figure 1 is an *Exotic Shorthair* or a *British Shorthair*. However, with detailed descriptions and reference images of both breeds, we can easily identify the correct result. Based on these observations, we propose **Prova**, a multi-modal **prototypes** classifier for **vast-vocabulary** object detection. Specifically, we utilize the detailed descriptions from the V3Det Challenge 2024 (Wang et al. 2024) to generate description embedding via a text encoder. We refer to these description embeddings as textual prototypes to replace the original class name embedding in previous alignment classifiers. Additionally, we collect images from the training dataset and exemplars in V3Det to generate visual prototypes via a visual encoder. Finally, we utilize the textual and visual prototypes for classification through dot product alignment with detected region features.

As a proof-of-concept, shown in Figure 2, we visualize the feature space of different prototypes. It illustrates that more detailed text descriptions enable the identification of more accurate semantic points within the feature space of the language modality, thereby enhancing the differentiation between categories. In addition, the feature space of the visual modality demonstrates greater distinction and reduced ambiguity compared to the language modality. By integrating both modalities, the classifier can better align with the detected objects, leading to improved object recognition results. As quantitative proof, we conduct experiments in supervised and open-vocabulary settings on V3Det. The results

show that our Prova is compatible with various detection frameworks, significantly improving object recognition accuracy in both settings with only two additional projection layers and four additional matrix multiplications. Furthermore, our best model with a **base-size** backbone, achieves a new state-of-the-art performance in both settings, particularly being $5\times$ lighter than the previous SoTA models with a **huge-size** backbone in the supervised setting. Our contributions can be summarized as:

- We point out that previous open-world classifiers using class names aligned with visual features are struggling when categories increase, and more detailed descriptions and reference images can alleviate this problem.
- We propose a simple yet novel approach, Prova, to extract multi-modal prototypes and leverage them to perform vast-vocabulary recognition, which is applicable to various detection frameworks with minimal modification.
- Our method demonstrates efficient and powerful results on vast-vocabulary detection datasets, e.g., V3Det, by only two additional projection layers and four additional matrix multiplications.

Related Work

Supervised Object Detection

Supervised object detection models, in a closed-set setting where the category vocabularies remain constant during both training and inference, can be divided into three frameworks: two-stage, one-stage, and DETR-based. Two-stage detectors (Ren et al. 2015; Cai and Vasconcelos 2018; Zhou, Koltun, and Krähenbühl 2021; Lin et al. 2017) first utilize a Region Proposal Network (RPN) to predict class-agnostic proposals that are pooled to region-of-interest features, and then recognize them by a classifier. One-stage detectors (Tian et al. 2022; Redmon and Farhadi 2017; Liu et al. 2016) directly regress and classify the anchor boxes or perform dense searches for visual cues. DETR-based detectors (Carion et al.

2020; Zhu et al. 2021; Zhang et al. 2023a; Meng et al. 2023a) treat object detection as a set prediction problem by utilizing an encoder-decoder transformer architecture.

Open-Vocabulary Object Detection

Leveraging the generalization capabilities of Vision-Language Models (Radford et al. 2021; Jia et al. 2021), open-vocabulary detectors are able to detect novel classes which are unseen during training. ViLD (Gu et al. 2022) effectively distills the knowledge of CLIP into a closed-set detector. MDETR (Kamath et al. 2021) and GLIP (Li et al. 2022) convert detection task as a grounding task to expand the training dataset. RegionCLIP (Zhong et al. 2022) uses image-caption data to construct region-wise pseudo-labels as extra training data. Detic (Zhou et al. 2022) makes use of image classification data to provide weak supervision. F-VLM (Kuo et al. 2023) directly builds a two-stage detector upon frozen CLIP and trains the detection heads only, and CORA (Wu et al. 2023b) adapts CLIP by region prompting and anchor pre-matching.

Prototype-based Object Recognition

Prototype is early studied in few-shot and zero-shot learning (Snell, Swersky, and Zemel 2017), and is also widely used in other fields (Xing et al. 2022; Zhang et al. 2023b; Yang et al. 2024; Meng et al. 2024). In object detection, the vision-language alignment classifiers can be treated as prototypes-based classifier with cosine distance. Besides that, FSOD^{up} (Wu et al. 2021) learns universal prototypes as condition to enhance the object classifier. GPA (Xu et al. 2020) propose a graph-induced prototype alignment framework for cross-domain detection. RPA (Zhang, Wang, and Mao 2021) separately aligns foreground and background RPN features by RPN prototype alignment method. In particular, MM-OVOD (Kaul, Xie, and Zisserman 2023) employs a method similar to ours by using a multi-modal classifier to enhance open-vocabulary performance. However, unlike MM-OVOD, we do not train additional modules. Instead, we analyze recognition problems from the perspective of vast-vocabulary detection and prototype learning, thereby enhancing the performance in both supervised and open-vocabulary settings.

Method

Preliminaries

Object detection can be conceptually divided into two parts: object localization and object classification. The former identifies the position of the object within the image by generating a bounding box, while the latter assigns a category label to the localized object.

Traditional detection models (Ren et al. 2015; Tian et al. 2022; Carion et al. 2020; Zhu et al. 2021) usually obtain object features from the entire image according to locations (Ren et al. 2015; Lin et al. 2017; Tian et al. 2022) or cross-attention mechanism (Carion et al. 2020; Zhu et al. 2021; Zhang et al. 2023a). Then, the object features are fed into a classification branch, *e.g.* MLP, to classify the region into predefined categories, and a location branch to compute

the bounding box, separately. Formally, we denote the detector as D , the classification branch as C , and the localization branch as L . Then, these processes can be expressed as:

$$X = D(I) \quad (1)$$

$$B = \mathcal{L}(X), \quad S = C(X) \quad (2)$$

where X denotes object features, B and S denotes bounding box predictions and classification logits, separately.

Model Overview

Our detection model primarily comprises a detector, a bounding box (Bbox) prediction head, and a classifier. As illustrated in Figure 3 (a), the image is processed by a detector to extract object features. Subsequently, these object features are utilized by the Bbox head to predict the coordinates of the Bbox and the multi-modal classifier to generate the corresponding category predictions. To enable the detection model to handle vast vocabularies in real-world scenarios, we utilize prior knowledge from both the language and visual modalities to initialize two additional alignment classifiers, and ensemble them as an entire multi-modal prototype classifier by averaging their logits.

Prototype Extraction

Textual Prototypes. As shown in Figure 2 (a), the textual prototypes generated by adding simple category names to fixed templates possess weak semantics and often fail to maintain a mapping relation with the vast image categories. To address this issue, as illustrated in Figure 3 (b), we utilize more detailed descriptive texts for each category from the V3Det Challenge (Wang et al. 2024), which were obtained by prompting GPT-4V (OpenAI 2023). These descriptions are encoded by a text encoder to produce more semantically rich embedding. We denote the descriptions as \mathcal{T} and text encoder as \mathbf{Enc}_t . The description embedding T can be extracted by:

$$T = \mathbf{Enc}_t(\mathcal{T}) : \mathbb{R}^{C \times L_t} \quad (3)$$

where C denotes the number of categories and L_t represents the latent dimension of description embedding. As shown in Figure 2 (b), detailed descriptions create a cleaner feature space compared to single vocabulary, and we use these description embedding as textual prototypes to guide the classification of the detection model.

Visual Prototypes. As shown in Figure 3 (b), we acquire reference images from example images of each category in V3Det, as well as images cropped from the ground-truth in the training set. For each category c ,

$$V^c = \mathbf{Enc}_v(\mathcal{I}_r^c) : \mathbb{R}^{n \times L_v} \quad (4)$$

where n represents the number of reference images for this category, and L_v represents the latent dimension of image embedding. We then perform weighted fusion based on the image resolution size which can be formally expressed by:

$$V = \mathcal{W} \cdot V^c : \mathbb{R}^{C \times n} \times \mathbb{R}^{n \times L_v} \rightarrow \mathbb{R}^{C \times L_v} \quad (5)$$

where \mathcal{W} denotes weighted matrix based on the image resolution size, C denotes the number of categories and L_v represents the latent dimension of image embedding.

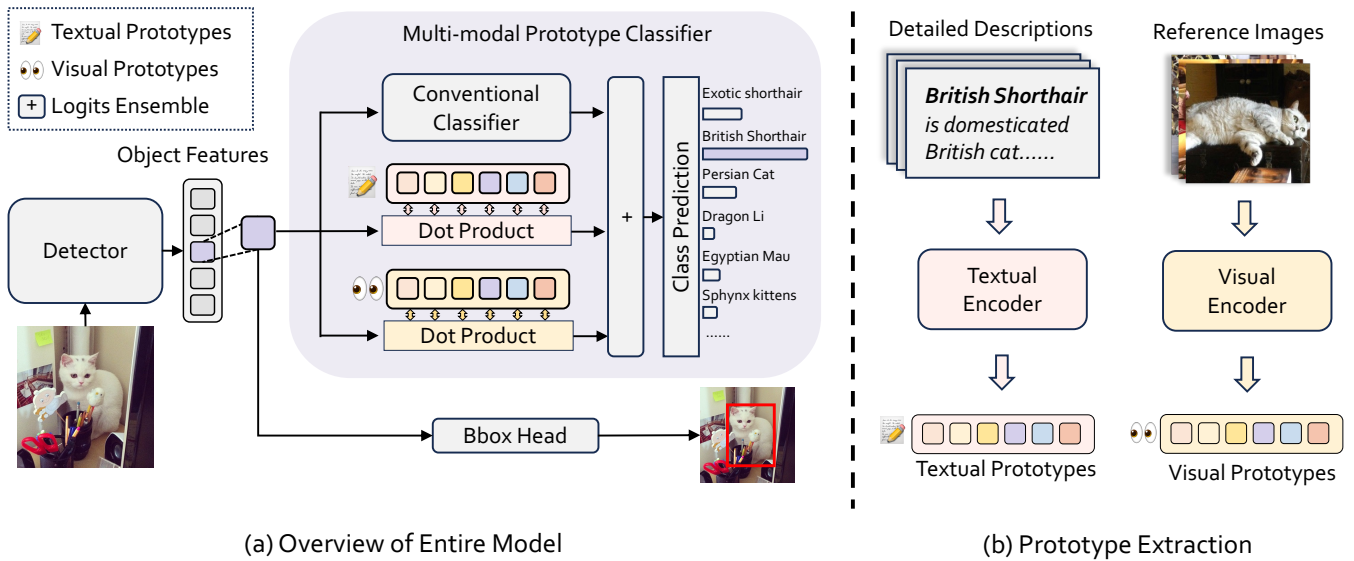


Figure 3: (a) Overview of Entire Model. Images are processed by a detector, *e.g.* Faster R-CNN (Ren et al. 2015), FCOS (Tian et al. 2022) or DINO (Zhang et al. 2023a) to extract object features (RoI features for Faster R-CNN, class features for FCOS and object queries for DINO), then these object features are utilized by the bounding box prediction head to predict bounding boxes and the Multi-modal Prototypes Classifier to generate corresponding category predictions. (b) Prototype Extraction. Textual prototypes are encoded by the LongCLIP-Large (Zhang et al. 2024) text encoder using detailed descriptions from the V3Det Challenge (Wang et al. 2024), and visual prototypes are encoded by the CLIP-ViT-Large (Radford et al. 2021) visual encoder using images from V3Det (Wang et al. 2023) training dataset and examples.

We utilize CLIP-ViT-Large as our visual encoder, and the embedding V as visual prototypes to guide the classification of the detection model. As shown in Figure 2 (c), the feature space of the visual modality exhibits greater distinction and less ambiguity compared to that of the language modality.

Multi-modal Prototype Classifier

Conventional Classifier. Let the object features be $X \in \mathbb{R}^{N \times L}$, where N represents the number of objects (*i.e.*, the number of anchors or queries) and L represents the latent dimension. The classifier can be expressed as follows:

$$S_{con} = W \cdot X : \mathbb{R}^{C \times L} \times \mathbb{R}^{L \times N} \rightarrow \mathbb{R}^{C \times N} \quad (6)$$

Where $S_{con} \in \mathbb{R}^{C \times N}$ denotes the classification logits, $W \in \mathbb{R}^{C \times L}$ denotes the learnable classifier parameters, and C denotes the number of categories.

We can also consider the conventional closed-set classifier as a specialized alignment method that learns abstract prototypes W and performs classification by calculating the similarity between W and X . Given that W lacks any prior knowledge and is initialized randomly, it converges slowly, exhibits weak generalization capability, and poses challenges when expanding to new categories.

Textual Prototype Classifier. Textual prototype classifier employs textual prototypes as prior knowledge to perform alignment and then classifies the object based on cosine similarity. Specifically, let the textual prototype be $T \in \mathbb{R}^{C \times L_t}$, C denotes the number of categories and L_t represents the latent dimension of textual prototypes. The textual prototype classifier can be expressed as follows:

$$S_{text} = \tau \cdot \frac{T \cdot (P_t \cdot X)}{\|T\| \cdot \|P_t \cdot X\|} \quad (7)$$

Where $S_{text} \in \mathbb{R}^{C \times N}$ denotes the classification logits by textual prototype classifier, $P_t \in \mathbb{R}^{L_t \times L}$ represents the projection layer that maps the latent dimension of X to the latent dimension of T , τ denotes the temperature of logits scale.

Visual Prototype Classifier. Similar to the textual prototype classifier, the visual prototype classifier also employs alignment method to enhance classification performance by utilizing visual prototypes. Specifically, let the visual prototype be $V \in \mathbb{R}^{C \times L_v}$, C denotes the number of categories and L_v represents the latent dimension of visual prototypes. The visual prototype classifier can be expressed as follows:

$$S_{vis} = \tau \cdot \frac{V \cdot (P_v \cdot X)}{\|V\| \cdot \|P_v \cdot X\|} \quad (8)$$

Where $S_{vis} \in \mathbb{R}^{C \times N}$ denotes the classification logits by visual prototype classifier, $P_v \in \mathbb{R}^{L_v \times L}$ represents the projection layer that maps the latent dimension of X to the latent dimension of V , τ also denotes the temperature of logits scale.

Formally, the prototype classifier provides semantically rich initialization matrices T and V to replace the role of W in the conventional classifier. As a result, it achieves better classification performance and faster convergence. Finally, we simply ensemble all classifiers for supervised setting, and

Method	Backbone	Epochs	AP	AP_{50}	AP_{75}
Faster R-CNN*	ResNet-50	24	25.4	32.9	28.1
Faster R-CNN* + Prova	ResNet-50	24	28.7 (+3.3)	37.1 (+4.2)	31.7 (+3.6)
FCOS*	ResNet-50	24	9.4	11.7	10.1
FCOS* + Prova	ResNet-50	24	15.6 (+6.2)	19.1 (+7.4)	16.7 (+6.6)
Deformable DETR	ResNet-50	50	34.4	39.9	36.4
DINO	ResNet-50	24	33.5	37.7	35.0
DINO + Prova	ResNet-50	24	36.4 (+2.9)	41.3 (+3.6)	38.1 (+3.1)
Faster R-CNN*	Swin-Base	24	37.6	46.0	41.1
FCOS*	Swin-Base	24	21.0	24.8	22.3
Deformable DETR	Swin-Base	50	42.5	48.3	44.7
DINO	Swin-Base	24	42.0	46.8	43.9
DINO + Prova	Swin-Base	24	44.5 (+2.5)	49.9 (+3.1)	46.6 (+2.7)

Table 1: Main results of supervised object detection on V3Det validation set. All the backbones are pre-trained on ImageNet-1K (Deng et al. 2009). * indicates utilizing Norm Linear Layer in the classification head (Wang et al. 2021).

Method	Backbone	Params	$\mathcal{D}^{backbone}$	$\mathcal{D}^{detector}$	AP	AP_{50}	AP_{75}
Cascade Mask R-CNN	EVA-ViT-G	1.1B	Merged-30M	Object365	49.4	54.8	51.4
DINO + Prova	Swin-Base	213.4M	ImageNet-22K	×	50.3	56.1	52.6

Table 2: Final results of supervised object detection on V3Det validation set. Params denotes the size of entire models. $\mathcal{D}^{backbone}$ indicates the pre-training dataset of backbones and $\mathcal{D}^{detector}$ indicates the pre-training dataset of detectors. Merged-30M (Fang et al. 2023) is a mixed dataset including ImageNet-22K, CC12M, CC3M, Object365, COCO and ADE20K.

remove closed-set classifier for open-vocabulary setting:

$$S = \begin{cases} \frac{S_{con} + S_{text} + S_{vis}}{3} & \text{if supervised} \\ \frac{S_{text} + S_{vis}}{2} & \text{else} \end{cases} \quad (9)$$

Experiments

In this section, we evaluate our method with various baselines on V3Det validation set in both supervised and open-vocabulary manners. We conduct additional experiments on LVIS in an open-vocabulary setting. Finally, we present ablation studies on different components of our classifier.

Supervised Object Detection

Datasets. We conduct experiments on V3Det, a vast-vocabulary object detection dataset containing 13204 categories with \sim 200K images. We train our models with all 13204 categories in V3Det train set with \sim 180K images, and evaluate them on the validation set with \sim 30K images.

Baselines. We consider the following baselines: (1) Faster R-CNN (Ren et al. 2015): a two-stage object detector. (2) FCOS (Tian et al. 2022): a one-stage object detector. (3) DINO (Zhang et al. 2023a): a DETR-based model. These baselines are to validate our proposed Prova. Moreover, we compare our best model with ImageNet-22k pre-trained Swin-Base (Liu et al. 2021) as the backbone to the previous SoTA detector with EVA-Huge (Fang et al. 2023).

Implementation Details. We use $8 \times V100$ for all experiments. For Faster R-CNN and FCOS, we follow the setups from V3Det without any modification except halving the batch size of FCOS. For DINO, we set the initial learning rate as 0.0002 with 24 epochs and multiplied 0.1 at the 20th epoch. We obtain descriptions from V3Det Challenge (Wang et al. 2024) and collect reference images from example images and training set in V3Det no more than 100 per category. The object features used to recognize objects are RoI features for Faster R-CNN, dense grid features for FCOS, and object queries for DINO. Following V3Det (Wang et al. 2023), we adopt repeat factor sampler (Gupta, Dollar, and Girshick 2019) and report average precision under different IoU noted as AP , AP_{50} , AP_{75} .

Results. Table 1 shows the result on V3Det val set for supervised setting, and all the backbones pre-trained on ImageNet-1K (Deng et al. 2009) following V3Det. As shown in this table, when using a ResNet-50 (He et al. 2016) as the backbone, different frameworks consistently benefit from our proposed Prova (*e.g.*, +3.3 AP on Faster R-CNN, +6.2 AP on FCOS and +2.9 AP on DINO). These results demonstrate that Prova is a general method applicable to popular detection frameworks. When using a Swin-Base (Liu et al. 2021) as the backbone, our method largely outperforms the baselines (*e.g.*, +2.5 AP on DINO), suggesting that Prova is also an efficient method and suitable for larger models.

Moreover, we compare our biggest model, DINO + Prova with Swin-Base pre-trained on ImageNet-22K (Deng et al. 2009), to a strong Cascade Mask R-CNN detector (Cai

Method	Backbone	Extra Training Data	Epochs	AP_{base}	AP_{novel}	AP_{final}
DINO*	ResNet-50	×	24	19.0	1.9	10.5
RegionCLIP	ResNet-50	×	48	22.1	3.1	12.6
CenterNet2*	ResNet-50	×	48	28.6	3.0	15.8
DST-Det	ResNet-50	×	24	-	7.2	-
DINO + Prova	ResNet-50	×	24	31.4	9.5	20.5
Detic	ResNet-50	ImageNet-22K	48	30.2	6.7	18.5
DINO + Prova	ResNet-50	ImageNet-22K	24	32.8	11.0	21.9

Table 3: Results of open-vocabulary object detection on V3Det validation set. * indicates replacing closed-set classifier with class name alignment classifier; - indicates that do not have reported number. AP_{final} is the mean of AP_{base} and AP_{novel} .

Method	Backbone	Backbone Size	$D^{backbone}$	V-L Pretraining	AP_r	AP
RegionCLIP	ResNet-50x4	87M	CLIP-400M	✓	22.0*	32.3
CORA	ResNet-50x4	87M	CLIP-400M	✓	22.2	-
CLIPSelf	ViT-B/16	86M	Merged-2B	✓	25.3*	-
RO-ViT	ViT-B/16	86M	ALIGN	✓	28.4	31.9
CFM-ViT	ViT-B/16	86M	ALIGN	✓	29.6	33.8
DINO + Prova	Swin-Base	88M	ImageNet-22K	×	31.5	44.2

Table 4: Results of open-vocabulary object detection on OV-LVIS. $D^{backbone}$ indicates pre-training data of backbone; V-L Pre-training indicates whether using contrastive image-text pre-training; * indicates reporting mask AP.

and Vasconcelos 2018), which uses EVA-ViT-G (Fang et al. 2023; Dosovitskiy et al. 2020) as backbone network pre-trained on Merged-30M, which includes ImageNet-22K, CC12M (Changpinyo et al. 2021), CC3M (Sharma et al. 2018), Object365 (Shao et al. 2019), COCO (Lin et al. 2014), and ADE20K (Zhou et al. 2017). Then the detector is pre-trained again on Object365. As shown in Table 2, despite the model being five times smaller and under a much weaker pre-training strategy, Prova outperforms this detector by 0.9 AP, demonstrating its efficiency and remarkable performance.

Open-vocabulary Object Detection

Datasets. We further experiments on V3Det under an open-vocabulary setting, where the classes are split into 6709 base classes and 6495 novel classes. We train our models on base classes with $\sim 130K$ images, and evaluate them on the validation set including both base classes and novel classes. We also utilize extra data from ImageNet-22K which overlaps the vocabularies of V3Det.

Baselines. We choose four different models: (1) RegionCLIP (Zhong et al. 2022) trained on base classes for 24 epochs to initialize the offline RPN and 24 epochs again to align regions and texts. (2) CenterNet2 (Zhou, Koltun, and Krähenbühl 2021) trained on V3Det base classes. (3) DST-Det (Xu et al. 2024) trained on V3Det base classes. (4) Detic (Zhou et al. 2022) trained on both V3Det base classes and ImageNet-22K overlapped classes. The backbones of four models are pre-trained on ImageNet-22k. We also report DINO with class name alignment classifier as the baseline.

Implementation detail. We utilize DINO with Prova as our open-vocabulary model. The backbone network is also

pre-trained on ImageNet-22k. For mixed dataset training, we follow the training recipe in RichSem (Meng et al. 2023b) to sample images from V3Det and ImageNet-22K in a 1:1 ratio; for data from ImageNet-22k, we treat a whole-image box as ground-truth bounding box and group four images together to build a mosaic following (Meng et al. 2023b). Again, we set the initial learning rate as 0.0002 with 24 epochs and multiply 0.1 at the 20-th epoch. The textual prototypes and base class visual prototypes are the same as ones in the supervised setting. In contrast, for novel classes, we obtain reference images from example images in V3Det Challenge (Wang et al. 2024) 1 image per category and ImageNet-22K 100 images per category if overlapped, to ensure that no novel images from the V3Det training set are seen.

Results. In Table 3, Prova largely outperforms the open-vocabulary models whether or not to use extra training data. Compared to the strong baseline models, *i.e.* CenterNet2/Detic (Zhou et al. 2022; Wang et al. 2023), Prova obtain $+4.3 AP_{novel}$ and $+3.4 AP_{final}$ when using extra data, and $+6.5 AP_{novel}$ and $+4.7 AP_{final}$ without extra data, demonstrating that our model gains strong generalization capability.

Additional Dataset. We evaluate Prova on LVIS (Gupta, Dollar, and Girshick 2019) which has 1203 categories. We simply extract visual prototypes by less than 100 images cropped from the LVIS training-set, and directly treat the class-name embedding as the textual prototypes. As shown in Table 4, unlike other detectors (Zhong et al. 2022; Wu et al. 2023b,a; Kim, Angelova, and Kuo 2023b,a) utilizing vision-language pretraining into model design, our offline prototype-based method effectively achieves a new state-of-the-art performance under the base-size comparison.

Class-name Alignment	Description Alignment	Visual Alignment	Conventional Classifier	AP	AP_{50}	AP_{75}
✓				10.8	12.9	11.5
	✓			16.0	18.7	16.8
		✓		28.1	32.2	29.5
			✓	24.5	28.2	25.7
	✓	✓		29.4	33.6	30.9
	✓	✓	✓	34.0	38.8	35.6

Table 5: Ablation on different classifiers. If utilizing multiple classifiers, we ensemble them by averaging the prediction logits.

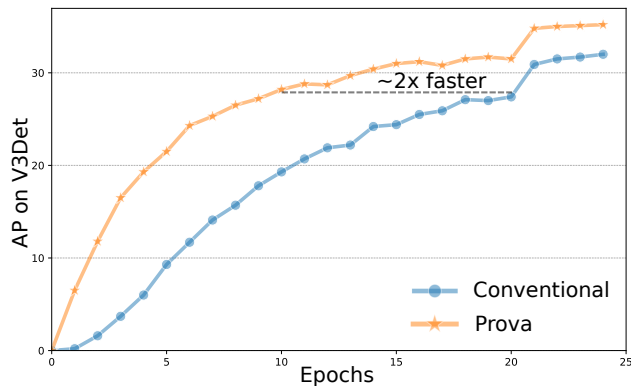


Figure 4: Convergence curves of DINO with conventional classifier and Prova for 24 epochs on V3Det val set.

Ablation Study

The evaluations are performed on V3Det dataset under supervised training of 12 epochs with DINO as detector and ResNet-50 pre-trained on ImageNet-1K as backbone.

Effectiveness of Different Classifier. We evaluate the effects of each component within the multi-modal prototype classifier. As shown in Table 5, compared to coarse class names, detailed descriptions enhance the differentiation between categories, and visual prototypes have greater distinction and reduced ambiguity. Integrating both prototypes, the classifier can better align with the detected objects.

Additionally, we plot the AP curves of Prova and conventional closed-set classifier trained for 24 epochs in a supervised setting. As show in Figure 4, Prova has faster convergence with multi-model prototypes compared to the conventional classifier, suggesting the efficiency of our method.

Comparisons with Different Encoders of prototypes. As shown in Table 6, we explore the visual prototype encoder with CLIP-ViT-Large, SigLIP-Large (Zhai et al. 2023) and DINOv2-Large (Oquab et al. 2024), finding that better results are achieved when utilizing CLIP, so we adopt it as our default visual encoder. In Addition, we also evaluate the textual prototype encoder with LongCLIP-Large, CLIP-ViT-Large and T5 (Raffel et al. 2020). As shown in Table 7, LongCLIP achieves better performance owing to its compatibility with long text input and training objective.

Effectiveness of the Number of Reference Images. We study the impact of visual prototype quality on the recognition results by using varying numbers of reference images. As shown in Table 8, the more reference images lead the better the result, suggesting that a greater number of reference images can result in more balanced visual prototypes.

Model	AP	AP_{50}	AP_{75}
CLIP-ViT-Large	28.1	32.2	29.5
SigLIP-Large	27.5	31.5	28.8
DINOv2-Large	25.5	29.4	26.8

Table 6: Ablation on visual prototype encoders.

Model	AP	AP_{50}	AP_{75}
LongCLIP-Large	16.0	18.7	16.8
CLIP-ViT-Large [†]	11.9	14.0	12.6
T5-3B	8.4	10.0	8.8

Table 7: Ablation on textual prototype encoders. †: Due to the maximum text token length, we truncate any part of a description that exceeds the limit.

Number of reference images	AP	AP_{50}	AP_{75}
1	21.3	24.4	22.2
10	26.6	30.5	27.9
100	28.1	32.2	29.5

Table 8: Ablation on the number of reference images.

Conclusion

We proposed **Prova**, a simple and effective multi-modal prototype classifier for vast-vocabulary object detection, to leverage multi-modal prototypes extracted by strong foundation models, alleviating the decline of object recognition performance when training categories scale up. Through extensive experiments, we demonstrated that our approach achieved state-of-the-art performance effectively.

Acknowledgements

This project was supported by NSFC under Grant No. 62102092.

References

- Cai, Z.; and Vasconcelos, N. 2018. Cascade r-cnn: Delving into high quality object detection. In *CVPR*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*.
- Changpinyo, S.; Sharma, P.; Ding, N.; and Soricut, R. 2021. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Fang, Y.; Wang, W.; Xie, B.; Sun, Q.; Wu, L.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2023. Eva: Exploring the limits of masked visual representation learning at scale. In *CVPR*.
- Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2022. Open-vocabulary Object Detection via Vision and Language Knowledge Distillation. In *ICLR*.
- Gupta, A.; Dollar, P.; and Girshick, R. 2019. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*.
- Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. Mdetr-modulated detection for end-to-end multi-modal understanding. In *ICCV*.
- Kaul, P.; Xie, W.; and Zisserman, A. 2023. Multi-modal classifiers for open-vocabulary object detection. In *ICML*.
- Kim, D.; Angelova, A.; and Kuo, W. 2023a. Contrastive feature masking open-vocabulary vision transformer. In *CVPR*.
- Kim, D.; Angelova, A.; and Kuo, W. 2023b. Region-aware pretraining for open-vocabulary object detection with vision transformers. In *CVPR*.
- Kuo, W.; Cui, Y.; Gu, X.; Piergiovanni, A.; and Angelova, A. 2023. Open-Vocabulary Object Detection upon Frozen Vision and Language Models. In *ICLR*.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022. Grounded language-image pre-training. In *CVPR*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *ICCV*.
- 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 *ECCV*.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *ECCV*.
- 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 *ICCV*.
- Meng, L.; Dai, X.; Chen, Y.; Zhang, P.; Chen, D.; Liu, M.; Wang, J.; Wu, Z.; Yuan, L.; and Jiang, Y.-G. 2023a. Detection hub: Unifying object detection datasets via query adaptation on language embedding. In *CVPR*.
- Meng, L.; Dai, X.; Yang, J.; Chen, D.; Chen, Y.; Liu, M.; Chen, Y.-L.; Wu, Z.; Yuan, L.; and Jiang, Y.-G. 2023b. Learning from rich semantics and coarse locations for long-tailed object detection. *NeurIPS*.
- Meng, L.; Lan, S.; Li, H.; Alvarez, J. M.; Wu, Z.; and Jiang, Y.-G. 2024. Segic: Unleashing the emergent correspondence for in-context segmentation. In *ECCV*.
- OpenAI. 2023. GPT-4V(ision) System Card. <https://api.semanticscholar.org/CorpusID:263218031>.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; HAZIZA, D.; Massa, F.; El-Nouby, A.; et al. 2024. DINOv2: Learning Robust Visual Features without Supervision. *TMLR*.
- 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 *ICML*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.
- Redmon, J.; and Farhadi, A. 2017. YOLO9000: better, faster, stronger. In *CVPR*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*.
- 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 *ICCV*.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *NeurIPS*.
- Tian, Z.; Chu, X.; Wang, X.; Wei, X.; and Shen, C. 2022. Fully convolutional one-stage 3d object detection on lidar range images. In *NeurIPS*.
- van der Maaten, L.; and Hinton, G. 2008. Visualizing Data using t-SNE. *JMLR*.
- Wang, J.; Zang, Y.; Zhang, P.; Chu, T.; Cao, Y.; Sun, Z.; Liu, Z.; Dong, X.; Wu, T.; Lin, D.; et al. 2024. V3Det Challenge 2024 on Vast Vocabulary and Open Vocabulary

Object Detection: Methods and Results. *arXiv preprint arXiv:2406.11739*.

Wang, J.; Zhang, P.; Chu, T.; Cao, Y.; Zhou, Y.; Wu, T.; Wang, B.; He, C.; and Lin, D. 2023. V3Det: Vast Vocabulary Visual Detection Dataset. In *ICCV*.

Wang, J.; Zhang, W.; Zang, Y.; Cao, Y.; Pang, J.; Gong, T.; Chen, K.; Liu, Z.; Loy, C. C.; and Lin, D. 2021. Seesaw loss for long-tailed instance segmentation. In *CVPR*.

Wu, A.; Han, Y.; Zhu, L.; and Yang, Y. 2021. Universal-prototype enhancing for few-shot object detection. In *ICCV*.

Wu, S.; Zhang, W.; Xu, L.; Jin, S.; Li, X.; Liu, W.; and Loy, C. C. 2023a. CLIPSelf: Vision Transformer Distills Itself for Open-Vocabulary Dense Prediction. In *ICLR*.

Wu, X.; Zhu, F.; Zhao, R.; and Li, H. 2023b. Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In *CVPR*.

Xing, Z.; Li, H.; Wu, Z.; and Jiang, Y.-G. 2022. Semi-supervised single-view 3d reconstruction via prototype shape priors. In *ECCV*.

Xu, M.; Wang, H.; Ni, B.; Tian, Q.; and Zhang, W. 2020. Cross-domain detection via graph-induced prototype alignment. In *CVPR*.

Xu, S.; Li, X.; Wu, S.; Zhang, W.; Tong, Y.; and Loy, C. C. 2024. DST-Det: Simple Dynamic Self-Training for Open-Vocabulary Object Detection. *arXiv:2310.01393*.

Yang, Y.; Ma, C.; Ju, C.; Zhang, F.; Yao, J.; Zhang, Y.; and Wang, Y. 2024. Multi-modal Prototypes for Open-World Semantic Segmentation. *IJCV*.

Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *ICCV*.

Zhang, B.; Zhang, P.; Dong, X.; Zang, Y.; and Wang, J. 2024. Long-CLIP: Unlocking the Long-Text Capability of CLIP. *arXiv:2403.15378*.

Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L.; and Shum, H.-Y. 2023a. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. In *ICLR*.

Zhang, H.; Wu, Z.; Wang, Z.; Chen, Z.; and Jiang, Y.-G. 2023b. Prototypical residual networks for anomaly detection and localization. In *CVPR*.

Zhang, Y.; Wang, Z.; and Mao, Y. 2021. Rpn prototype alignment for domain adaptive object detector. In *CVPR*.

Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Regionclip: Region-based language-image pretraining. In *CVPR*.

Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene Parsing through ADE20K Dataset. In *CVPR*.

Zhou, X.; Girdhar, R.; Joulin, A.; Krähenbühl, P.; and Misra, I. 2022. Detecting twenty-thousand classes using image-level supervision. In *ECCV*.

Zhou, X.; Koltun, V.; and Krähenbühl, P. 2021. Probabilistic two-stage detection. *arXiv:2103.07461*.

Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2021. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *ICLR*.