

# Unbiased Region-Language Alignment for Open-Vocabulary Dense Prediction

Yunheng Li<sup>1</sup>, Yuxuan Li<sup>1</sup>, Quan-Sheng Zeng<sup>1</sup>, Wenhai Wang<sup>3,4</sup>, Qibin Hou<sup>1,2,†</sup>, Ming-Ming Cheng<sup>1,2</sup>

<sup>1</sup>VCIP, CS, Nankai University      <sup>2</sup>NKIARI, Shenzhen Futian

<sup>3</sup>OpenGVLab, Shanghai AI Laboratory      <sup>4</sup>The Chinese University of Hong Kong

[yunhengli@mail.nankai.edu.cn](mailto:yunhengli@mail.nankai.edu.cn), [yuxuan.li.17@ucl.ac.uk](mailto:yuxuan.li.17@ucl.ac.uk)

## Abstract

Pre-trained vision-language models (VLMs), such as CLIP, have demonstrated impressive zero-shot recognition capability, but still underperform in dense prediction tasks. Self-distillation recently is emerging as a promising approach for fine-tuning VLMs to better adapt to local regions without requiring extensive annotations. However, previous state-of-the-art approaches often suffer from significant ‘foreground bias’, where models tend to wrongly identify background regions as foreground objects. To alleviate this issue, we propose DenseVLM, a framework designed to learn unbiased region-language alignment from powerful pre-trained VLM representations. DenseVLM leverages the pre-trained VLM to retrieve categories for unlabeled regions and then decouples the interference between foreground and background features. This separation ensures accurate region-category alignment while maintaining semantic distinctions during training. We show that DenseVLM can directly replace the original VLM in open-vocabulary object detection and image segmentation methods, leading to notable performance improvements. Furthermore, it exhibits promising zero-shot scalability when training on more extensive and diverse datasets. Our code is publicly available <https://github.com/HVision-NKU/DenseVLM>.

## 1. Introduction

Open-vocabulary dense prediction, primarily including object detection [13, 20, 27, 46, 66] and image segmentation [5, 12, 33, 36, 61, 62], aims to identify regions or dense visual concepts of arbitrary categories as described by the text. Benefiting from the powerful pre-trained Vision-Language Models (VLMs), recent open-vocabulary approaches [25, 62, 72] for dense prediction have achieved great progress.

Popular VLMs, such as CLIP [50] and EVA-CLIP [53],

<sup>†</sup>Corresponding Author. Email: [houbq@nankai.edu.cn](mailto:houbq@nankai.edu.cn)

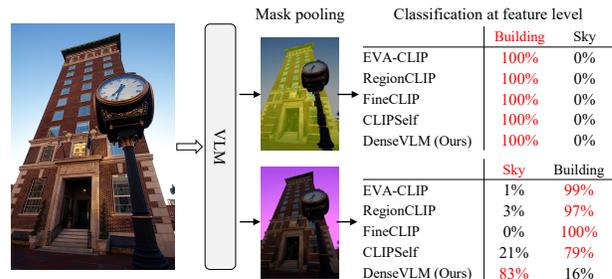


Figure 1. Illustration of foreground bias. Previous methods [53, 57, 69] often produce similar foreground predictions for background regions, our approach effectively alleviates this issue.

have exhibited remarkable zero-shot recognition abilities for global image understanding [11, 17, 18, 34, 35]. However, these models expose notable limitations in the understanding of local visual semantics, particularly in localizing and identifying small objects and background stuff [69, 72]. This limitation arises from the training manner of VLMs that align images with global text while neglecting the correspondences between local image regions and their specific text descriptions. To alleviate this issue, some studies use region-text or pseudo region-text pairs [32, 39, 68, 69] but these methods are limited by the high annotation cost and lack scalability for open-world scenes. In contrast, self-supervised approaches, such as CLIPSelf [57] and MaskEmbed [6], align region semantics using cropped image representations or reconstruct masked patch embeddings, respectively. These self-distillation approaches, which do not rely on annotated data, offer flexibility and scalability across a variety of datasets.

Despite the advances, previous VLMs [50, 53, 69] pre-trained on image-text and region-text pairs tend to overemphasize the foreground objects at the expense of the background context. This disproportionate focus results in a pronounced predilection for foreground object recognition and causes models to incorrectly associate background regions with foreground labels. Consequently, in dense prediction tasks, these VLMs often misclassify background regions to co-occurring foreground classes—a phenomenon we term as ‘foreground bias.’ To illustrate this issue, we present

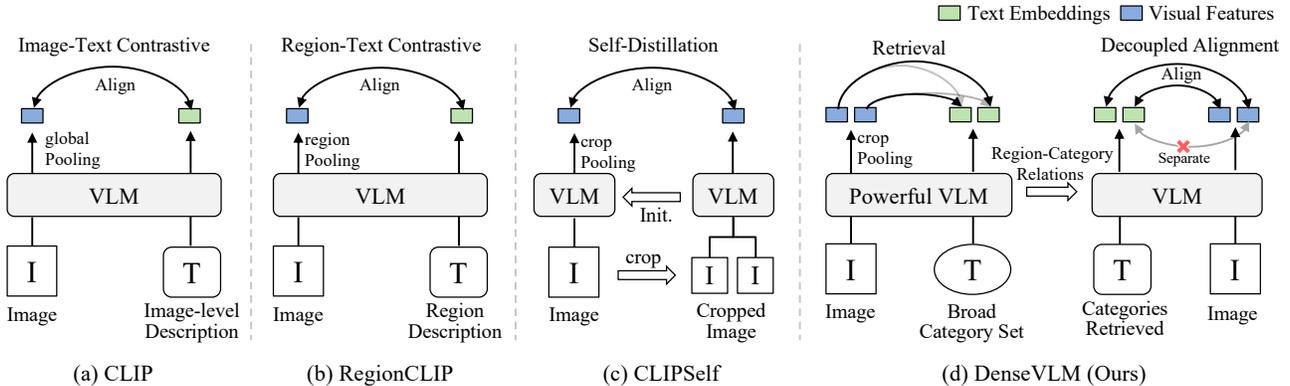


Figure 2. Comparison of different VLMs. Unlike existing methods using (a) image-text contrastive learning [50], (b) region-text contrastive learning [69] or (c) self-distillation [57], our method leverages powerful model representations for region-language alignment.

an example comparing the classification results of popular VLMs [53, 57, 69], where features are extracted from regions using ground-truth masks. As illustrated in Fig. 1, these models tend to confuse ‘sky’ (a background class) with ‘building’ (a foreground object).

To address this bias, we propose aligning foreground and background regions separately, ensuring explicit semantic separation through distinct category sets. To achieve this, we introduce DenseVLM, an end-to-end framework designed for unbiased region-language alignment, as shown in Fig. 2(d). Specifically, for unlabeled regions, we leverage a powerful pre-trained VLM that has learned robust semantic features from diverse data to retrieve relevant categories without relying on paired data or self-distillation [53, 57, 69]. To ensure semantic diversity, DenseVLM incorporates a broad and comprehensive set of categories derived from large-scale datasets or predicted by generative models [65]. A key feature of DenseVLM is its ability to classify regions as either foreground or background based on predefined sets for each category. This classification way enables the decoupling of region features, reducing interactions between foreground and background. By decoupling these features, DenseVLM achieves independent region alignment of their respective categories while maintaining a semantic separation during training. Furthermore, DenseVLM improves both efficiency and performance by directly extracting region features from the dense features of VLMs, avoiding traditional image cropping [57, 73].

We evaluate the effectiveness of DenseVLM on several open-vocabulary benchmarks [57], including object detection (box classification) and image segmentation (thing and stuff mask recognition). DenseVLM is adaptable to various network architectures, such as ViTs [9] and CNNs [30], and consistently outperforms other competing methods [50, 53, 57]. Furthermore, DenseVLM has great scaling ability, showing promising performance improvement when scaling up the training set based on the SA-1B [26] dataset. For downstream tasks, DenseVLM improves the two-stage

models [27] of OV-COCO [3] and OV-LVIS [15] in open-vocabulary object detection and achieves notable gains in open-vocabulary semantic segmentation over state-of-the-art methods such as SAN [62] and Cat-Seg [5]. We summarize our contributions as follows.

- We identify the foreground bias issue in existing VLMs and propose region-text alignment by incorporating explicit semantic structuring through category guidance.
- We propose DenseVLM, a region-language alignment framework that leverages a powerful VLM to retrieve categories for unlabeled regions and decouples foreground and background features to reduce foreground bias.
- Extensive experiments on dense prediction benchmarks show that our DenseVLM outperforms previous methods and exhibits promising scalability.

## 2. Related work

**Open-vocabulary dense prediction.** Open-vocabulary dense prediction approaches aim to overcome the constraints of predefined categories, thereby enhancing their application in object detection [13, 27, 46, 66] and image segmentation [12, 31, 33, 36, 61, 62]. The success of pre-trained vision-language models like CLIP [50], has further inspired advancements in this area. In open-vocabulary detection, recent studies [13, 56] exploit the CLIP models to effectively identify novel objects. Furthermore, several works [27, 58] construct object detectors that utilize frozen CLIP encoders, thereby reducing computational overhead while maintaining performance. For open-vocabulary segmentation, a common two-stage pipeline [7, 8, 41, 61] integrates a class-agnostic mask generator with a fixed CLIP encoder to achieve cross-modal alignment and mask classification. Recent methods also investigate the use of shared frozen CLIP with side adapter networks [63] or the adoption of single-stage frameworks [64]. However, due to its training on image-text pairs, CLIP lacks precise local vision-language alignment, which is essential for dense prediction

tasks. Although recent studies fine-tune CLIP for pixel-level [5, 60] or mask-level [22] alignment, they are constrained by the scarcity of densely labeled data.

**Vision-language alignment at image and region levels.** Pre-training visual-language models enable alignment between image and text representations [2, 19, 24, 44, 45, 50, 55]. By using image-text pairs [52] as training data, methods such as CLIP [50] and ALIGN [19] exhibit impressive zero-shot classification capabilities. To improve the alignment of vision-language of dense features, some training-free studies [28, 29, 54, 59] attempt to modify the output layers of CLIP. Recent methods [48, 49, 51] align visual patches with image-level text or learn local-global correspondence via self-distillation, both requiring extensive text-image pairs.

For precise local region alignment, researchers utilize annotations from visual grounding datasets [39] to train models on region-text alignment. For example, some methods such as GLIP [32, 68] and Grounding DINO [39] learn region-language grounding from region-text pairs or generates pseudo region-text pairs like RegionCLIP [69]. In the context of open-vocabulary detection and segmentation, some works achieve dense visual and text alignment by using mask attention [8, 25, 62, 72] or fine-tuning CLIP [5, 21, 60]. However, these methods are constrained by the high cost of annotation, making large-scale deployment challenging. To address the issue of annotated data scarcity, recent methods like CLIPSelf [57] use cropped images for semantic alignment, while MaskEmbed [6] leverages masked patch embeddings for feature reconstruction. Despite these advances, the efficacy of self-distillation methods is constrained by the suboptimal performance of the teacher model and is further compromised by foreground bias. To overcome these limitations, we leverage a powerful VLM to retrieve categories for unlabeled regions while decoupling foreground and background features through textual category guidance.

### 3. Method

Our aim is to develop a region-language alignment model that can effectively align local visual and semantic features, thereby enhancing the performance of VLMs in open-vocabulary dense prediction tasks. To achieve this, it is crucial to alleviate the foreground bias problem prevalent in previous VLMs [50, 53] that arises from training on image-text pairs. Moreover, our approach seeks to receive better performance beyond the constraints of self-distillation [57].

#### 3.1. VLM’s representation

VLMs are typically designed to learn both global visual and textual representations within a shared semantic space. In ViT-based VLMs [4, 53], dense visual features are extracted through residual attention blocks. Following [57, 72], we

derive the dense image feature map  $\mathcal{F}$  by modifying the final residual attention block, *i.e.*, removing the SoftMax operation and incorporating mapping layers. For a set of categories  $\{c^1, \dots, c^D\}$ , where  $D$  is the total number of categories, textual descriptions are generated using a prompt-template strategy [13], like “*This is a photo of the  $c$  in the scene.*” These prompts are then encoded into text embeddings  $\mathcal{T} = \{\mathcal{T}^1, \dots, \mathcal{T}^D\}$  by the text encoder.

#### 3.2. DenseVLM framework

We propose DenseVLM, an end-to-end region-language alignment framework designed to mitigate foreground bias. DenseVLM achieves this by precisely aligning unlabeled regions with their corresponding categories. As shown in Fig. 3, the framework consists of two key components. First, it retrieves category semantics for region features using the P-VLMs (Powerful VLMs) with frozen weights. Second, it decouples these region-language alignments into foreground and background to train U-VLMs (Unbiased VLMs) without foreground bias. In particular, DenseVLM operates without requiring any annotations, relying on diverse category semantics from large-scale datasets [1, 71] or generating category sets from images using the NXP [65].

**Image patches to semantic regions.** To achieve region alignment, the patch-level visual features from Sec. 3.1 need to be transformed into semantic region features. We adopt a strategy similar to [57], dividing the dense feature map into an  $m \times n$  grid of patches. Unlike this approach, we refrain from directly cropping the input image, thereby enhancing both computational efficiency and representation effectiveness. In each training iteration,  $m$  and  $n$  are randomly selected from the set  $\{2, \dots, M\}$ , where  $M$  defaults to 6, allowing for varying patch grid sizes. The semantic region representation  $\{\mathcal{F}^1, \dots, \mathcal{F}^{m \times n}\}$  is then extracted from dense feature map  $\mathcal{F}$  via pooling (RoIAlign) [16]. This patch sampling strategy effectively captures region features of both foreground objects and background scenes. However, due to the foreground bias in VLMs [50, 53, 57], we observe that background regions are usually misclassified as foreground classes, despite containing only a small proportion of foreground-related patches.

**Powerful VLM for retrieval regions.** Building on the dense representation extracted from VLM described in Sec. 3.1 and its ability to map image patches to semantic regions, the powerful VLM enables the extraction of both region features  $\mathcal{F}_P = \{\mathcal{F}_P^1, \dots, \mathcal{F}_P^{m \times n}\}$  and text embeddings  $\mathcal{T}_P = \{\mathcal{T}_P^1, \dots, \mathcal{T}_P^D\}$ . Next, unlabeled regions are retrieved and matched with the most relevant categories by computing the cosine similarity between the region features  $\mathcal{F}_P$  and text embeddings  $\mathcal{T}_P$ . For a specific region  $k$ , the cosine similarity between its features  $\mathcal{F}_P^k = \mathcal{F}_P[k, \cdot]$  and

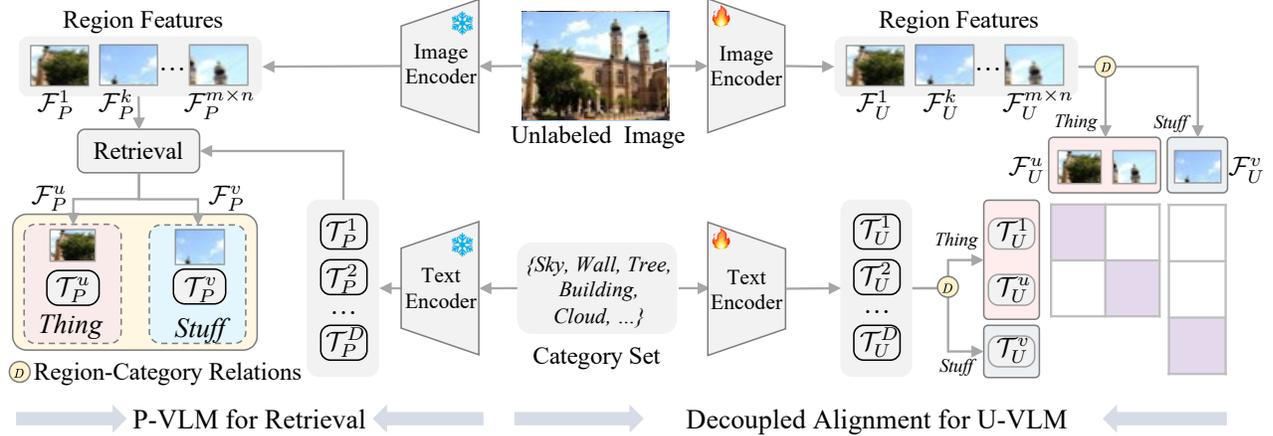


Figure 3. Overview of DenseVLM. DenseVLM leverages the powerful VLM to **retrieve** categories for unlabeled regions and distinguish between foreground and background. During VLM training, DenseVLM **decouples** interference between features of foreground and background regions, **aligning** each region with its corresponding text embeddings.

the text embeddings of all categories is computed as:

$$\cos(\mathcal{F}_P^k, \mathcal{T}_P^i) = \frac{\mathcal{F}_P^k \cdot \mathcal{T}_P^i}{\|\mathcal{F}_P^k\| \|\mathcal{T}_P^i\|}, \quad \forall i = 1, 2, \dots, D \quad (1)$$

where  $\cdot$  denotes the dot product and  $\|\cdot\|$  represents the Euclidean norm. The probability of associating this region with the categories is determined as follows:

$$p^k(y = c | \mathcal{F}_P^k, \mathcal{T}_P) = \frac{\exp(\cos(\mathcal{F}_P^k, \mathcal{T}_P^c) / \tau)}{\sum_{j=1}^D \exp(\cos(\mathcal{F}_P^k, \mathcal{T}_P^j) / \tau)}, \quad (2)$$

where  $\tau = 0.01$  is a temperature hyperparameter.

The use of random grids to extract region features often introduces uncertainty in fully covering a single object, especially in scenarios with multiple objects. This uncertainty significantly affects the precision of the region-category alignment. To mitigate this issue, we leverage a region denoising scheme by discarding any region whose matching probability falls below a threshold  $\theta$ , which is set to 0.3 by default. Therefore, this retrieval and denoising process yields a more precise and reliable alignment of regions with their corresponding categories. The optimal category for each region  $k$ , denoted as  $c^k$ , is determined by:  $c^k = \operatorname{argmax}(p^k)$ , where each region-category relation is represented as  $(k, c^k)$ . This alignment is fundamental for the subsequent decoupled alignment process.

**Decoupled region alignment to train VLMs without foreground bias.** Upon the established region-category relations, we perform alignment between the region and text representations to train the U-VLM. A straightforward approach would be to directly align region features and text embeddings of their corresponding categories while maximizing the separation of unrelated categories. However, since the U-VLM directly inherits from the P-VLM, our experimental results in Sec. 4.3 reveal that this method ex-

acerbates foreground bias, leading to improved foreground detection but limited gains in background recognition.

To mitigate this issue, we propose a decoupled alignment strategy that separates the alignment process for foreground and background regions. Specifically, we denote the region features of the training U-VLM as  $\mathcal{F}_U$ , using the same partitioning grid as the P-VLM. The text embeddings are represented as  $\mathcal{T}_U$ . By leveraging the region-category relations  $(k, c^k)$  retrieved by the P-VLM, we establish a one-to-one mapping for the U-VLM to associate the region features and their corresponding category embeddings. To distinguish the semantic regions of foreground and background, we decouple these region-category relations into two distinct groups, following the predefined category sets: foreground objects *Thing* ( $\mathcal{U}$ ) and background regions *Stuff* ( $\mathcal{V}$ ). By selectively contrasting against text embeddings related to irrelevant categories, we guide the model to focus more on the relevant background regions, reducing the impact of irrelevant foreground categories. This selective contrast helps the model capture the distinctive characteristics of background regions, leading to a more accurate separation between foreground and background. The alignment process can be effectively achieved by maximizing the cosine similarity for the region features and the text embeddings. According to Eqn. (2), when  $c^k \in \mathcal{V}$ , the probability  $q^k$  for a specific region is calculated as:

$$q^k = \frac{\exp(\cos(\mathcal{F}_U^k, \mathcal{T}_U^c) / \tau)}{\sum_{j=1}^{\mathcal{U} \cup \mathcal{V}} \exp(\cos(\mathcal{F}_U^k, \mathcal{T}_U^j) / \tau)}. \quad (3)$$

Similarly, when  $c^k \in \mathcal{U}$ , the probability  $\tilde{q}^k$  is computed as:

$$\tilde{q}^k = \begin{cases} \frac{\exp(\cos(\mathcal{F}_U^k, \mathcal{T}_U^c) / \tau)}{\sum_{j=1}^{\mathcal{U}} \exp(\cos(\mathcal{F}_U^k, \mathcal{T}_U^j) / \tau)} & \text{if } c \in \mathcal{U}, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

**End-to-end optimization.** As in Eqn. (3) and Eqn. (4),

Method	COCO						ADE20K					
	Boxes		Masks-T		Masks-S		Boxes		Masks-T		Masks-S	
	Top1	Top5										
OpenCLIP [4]	49.8	74.3	51.9	72.2	29.2	54.9	28.4	54.1	29.6	53.4	37.9	66.6
DFN [10]	38.3	65.0	31.0	57.0	26.4	54.9	30.6	57.9	24.2	49.9	32.2	57.7
SigLIP [67]	39.9	61.4	40.4	60.1	30.3	56.4	25.9	49.2	27.3	47.6	34.5	57.3
EVA-CLIP [53]	44.3	68.7	44.7	66.0	26.2	51.9	33.0	57.6	33.9	56.2	36.2	62.3
RegionCLIP <sup>†</sup> [69]	68.5	89.5	60.7	84.3	22.0	53.5	43.2	72.2	34.0	62.6	37.7	68.6
FineCLIP <sup>†</sup> [23]	64.7	86.1	62.5	80.9	36.9	70.3	43.9	71.2	45.5	68.6	46.0	74.8
CLIPSelf <sup>†</sup> [57]	69.1	88.2	66.7	83.0	41.7	75.2	48.1	77.7	47.5	74.2	53.7	82.8
DenseVLM <sup>†</sup> (Ours)	<b>72.3</b>	<b>89.9</b>	<b>70.1</b>	<b>84.4</b>	<b>44.9</b>	<b>76.4</b>	<b>51.0</b>	<b>81.8</b>	<b>49.3</b>	<b>76.5</b>	<b>57.0</b>	<b>84.0</b>

Table 1. Comparisons of different models on dense representation. We report the Top1 and Top5 mean accuracy on classifying boxes and panoptic masks (thing and stuff). <sup>†</sup> denotes models trained on the COCO and evaluated in a zero-shot setting on the ADE20K dataset.

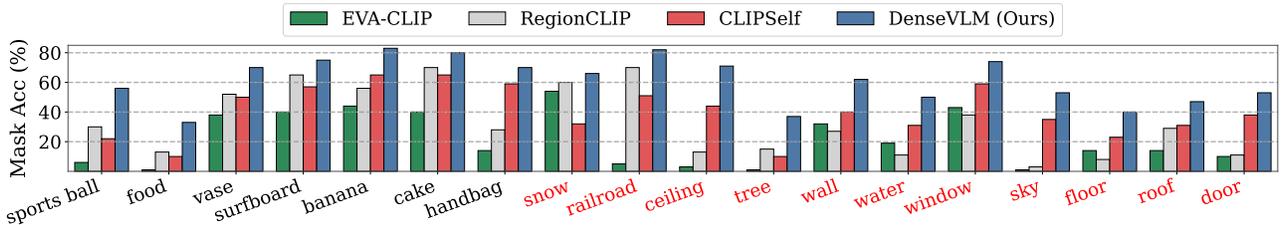


Figure 4. Mask accuracy comparison across categories in COCO dataset. Our method achieves notable improvements, especially in addressing foreground bias. The foreground categories are shown in black, and the background categories are highlighted in red.

we also compute  $p^k$  and  $\tilde{p}^k$  for P-VLM. The proposed method DenseVLM supports end-to-end training through KL-divergence optimization as follows.

$$\mathcal{L}^k = \begin{cases} \text{KL}(p^k || q^k) & \text{if } c^k \in \mathcal{U}, \\ \text{KL}(\tilde{p}^k || \tilde{q}^k) & \text{otherwise.} \end{cases} \quad (5)$$

The overall loss for each image is computed as  $\mathcal{L} = \frac{1}{m' \times n'} \sum_k^{m' \times n'} \mathcal{L}^k$ , where the sum is taken over all regions, excluding those removed by the region denoising scheme.

## 4. Experiments

### 4.1. Benchmarks

**Experiment settings.** To verify the effectiveness of the proposed DenseVLM, we perform experiments on dense prediction tasks using the COCO panoptic [38] val2017 split and ADE20K panoptic [70] val split. Following CLIPSelf [57], we evaluate box classification using pooled box features (**Boxes**) and mask classification with pooled mask features, distinguishing foreground objects (**Masks-T**) from background content (**Masks-S**). This process follows a similar procedure as illustrated in Fig. 1, where we use ground truth annotations to extract local features and evaluate the local classification. Results are reported in Top-1 and Top-5 mean accuracy across all experiments.

**Implementation details.** We employ the ViT-L/14 model from CLIPSelf [57] as the powerful P-VLM and ViT-B/16 from EVA-CLIP [53] as the training U-VLM. To en-

hance computational efficiency, the powerful VLM are kept frozen, and only the image encoder of the U-VLM is trained with pre-extracted text embeddings. Considering the practical application of the downstream tasks and the trade-off between performance and efficiency, we resize the input images to a uniform resolution of  $512 \times 512$  pixels. The models are trained for 6 epochs using the AdamW [43] optimizer with a weight decay of 0.1. The initial learning rate is set to  $1 \times 10^{-5}$  with a cosine decay [42] schedule.

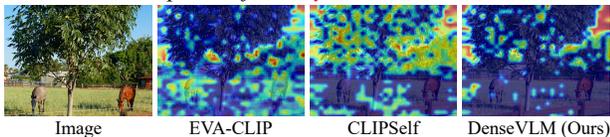
### 4.2. Comparison with other VLMs

**Quantitative evaluation.** We perform a comprehensive quantitative evaluation of dense representations across multiple VLMs on the COCO Panoptic [38] and ADE20K Panoptic [70] datasets. As shown in Tab. 1, while previous methods [4, 10, 53, 67] achieve strong zero-shot image classification performance, their performance in region recognition is notably suboptimal. For instance, EVA-CLIP [53] achieves only a Top-1 box classification accuracy of 44.3% on COCO and 33.0% on ADE20K. Although RegionCLIP [69], which is trained on region-text pairs, shows improved performance on COCO, but fails to generalize to datasets such as ADE20K. Moreover, FineCLIP [23] and CLIPSelf [57], which incorporate self-distillation, achieve higher Mask-T classification accuracy but perform poorly in Mask-S classification. In contrast, our method, DenseVLM, notably outperforms these models, achieving a 4.3% improvement in Top-1 Mask-T accuracy and a 3.9% improvement in Top-1 Mask-S accuracy on COCO, highlighting its

VLMs Frozen & Training	Region Cropping	Alignment Strategy	GPU Memory (per card)	Time Overhead (per epoch)	Boxes		Masks-T		Masks-S	
					Top1	Top5	Top1	Top5	Top1	Top5
ViT-B/16 & ViT-B/16	Images	Features KD	37G	25min	69.1	88.2	66.7	83.4	41.7	75.2
ViT-L/14 & ViT-B/16	Images	Features KD	39G	37min	24.2	52.4	23.4	51.1	10.1	39.1
ViT-L/14 & ViT-B/16	Images	Logics KD	39G	55min	72.2	89.8	68.8	83.8	42.6	75.2
ViT-L/14 & ViT-B/16	Features	DenseVLM	39G	<b>23min</b>	<b>73.4</b>	<b>90.5</b>	<b>71.0</b>	<b>84.8</b>	<b>45.6</b>	<b>77.8</b>

Table 2. Results and comparisons of various frameworks, including target VLMs for aligning representations, region cropping of frozen VLMs, and optimization strategies for region-language alignment, with GPU memory efficiency and time overhead. All models are trained on four A40 GPUs, with each epoch containing 118k images.

Text: “This is a photo of the *sky* in the scene.”



Text: “This is a photo of the *sofa* in the scene.”

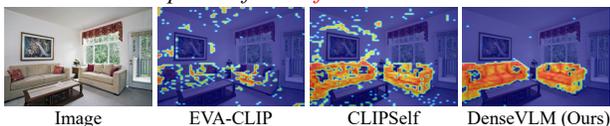


Figure 5. Comparing the alignment effect of our DenseVLM with other methods through visualizations of cosine similarity maps between visual features and text embeddings.

effectiveness in dense prediction tasks.

Fig. 4 shows a comparison of mask accuracy across different categories, further demonstrating the effectiveness of DenseVLM in region-text alignment. Notably, our method notably improves accuracy in background categories such as ‘sky’ and ‘wall’, which can mitigate foreground bias.

**Qualitative results.** We visualize attention maps using cosine similarity for text-described object categories. As shown in Fig. 5, DenseVLM achieves more accurate and complete object localization than EVA-CLIP and CLIPSelf. Moreover, it better preserves semantic separation, reducing interference from other objects.

### 4.3. Ablation study

In our ablation study, DenseVLM is trained on the unlabeled images from the COCO train2017 split using 4×A40 GPUs and evaluated on the val2017 split. The experiments in Sec. 4.3 utilize the ViT-B/16 model from EVA-CLIP [53] for its superior efficiency and capacity.

**Framework design exploration.** As shown in Tab. 2, we conduct ablation studies to explore key design choices in DenseVLM, including target VLMs for aligning representations, region cropping, and optimization strategies with a focus on GPU memory usage and time overhead. Using a self-distillation strategy [57] as the baseline, we observe a significant performance drop when replacing the target model with ViT-L/14 CLIP from CLIPSelf, due to disruption in visual-language alignment. Replacing feature distillation with logit distillation improves performance but de-

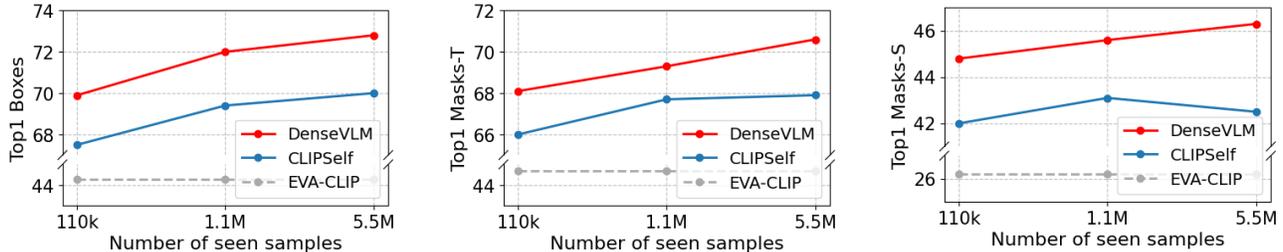
P-Thing		P-Stuff		Boxes	Masks-T	Masks-S
Thing	Stuff	Thing	Stuff			
✓	✓	✓	✓	<b>74.3</b>	70.9	42.6
✓	✗	✗	✓	74.2	70.8	42.0
✓	✓	✗	✓	74.1	70.9	41.3
✓	✗	✓	✓	73.4	<b>71.0</b>	<b>45.6</b>

Table 3. Ablation study on decoupled alignment. ✓ denotes that a region is separated from the categories in this set.

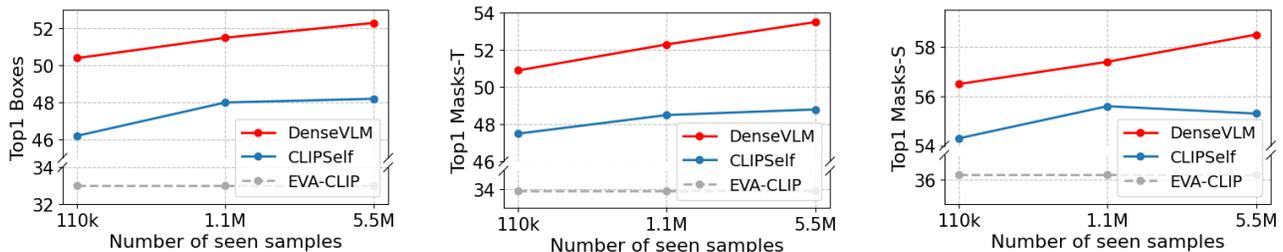
creases training efficiency due to repeated feature extraction while introducing foreground bias issues. In contrast, DenseVLM leverages a more efficient feature-cropping strategy and a decoupled alignment framework to achieve superior performance while substantially reducing both training time and GPU memory usage.

**Ablation study on decoupled alignment.** Tab. 3 analyzes the effect of the decoupled alignment strategy on DenseVLM. The retrieved category  $c^k$  can be categorized into *Thing* and *Stuff*, referred to as ‘P-Thing’ and ‘P-Stuff’, respectively. DenseVLM leverages  $c^k$  for regions to selectively contrast against irrelevant *Thing* and *Stuff* categories. When regions contrast against all categories, the model can better identify the foreground regions but struggles with distinguishing the background (1st row). In the fully decoupled setting, where each region contrasts against its own set of categories, performance further degrades (2nd row). The most pronounced foreground bias occurs when P-Thing contrasts against *Stuff* but P-Stuff does not contrast against *Thing* (3rd row). In DenseVLM, we adopt a balanced strategy where P-Stuff contrasts against both *Thing* and *Stuff*, while P-Thing does not contrast against *Stuff*, resulting in improved overall performance by mitigating bias (4th row).

**Ablation study on different category sets.** Tab. 4 presents the effects of different category sets used for retrieval and alignment performance. When using the category set from dataset annotations, we observe that as the number of categories increases, performance consistently improves. Adding only the background category helps the model distinguish foreground differences, as seen in rows 1 and 2, and incorporating more foreground categories leads to performance gains, as seen in rows 3 and 4. These improvements can be attributed to the enhanced representational capacity



(a) Top1 mean accuracy of models on COCO dense prediction benchmark.



(b) Top1 mean accuracy of models on ADE20K dense prediction benchmark.

Figure 6. Zero-shot comparisons of models pre-trained on datasets with three different scales. We select three training sets from the SA-1B dataset [26]: 100K, 1.1M, and 5.5M seen samples and perform the zero-shot evaluation on the COCO and ADE20K benchmarks.

Categories	Boxes		Masks-T		Masks-S	
	Top1	Top5	Top1	Top5	Top1	Top5
<i>The category set is from the dataset annotation.</i>						
133 (80)	71.1	88.5	68.7	83.0	44.7	75.2
171 (80)	72.3	89.8	69.4	85.8	44.2	76.0
273 (160)	72.3	89.9	70.1	84.4	44.9	76.4
316 (204)	<b>73.4</b>	<b>90.5</b>	<b>71.0</b>	<b>84.8</b>	<b>45.6</b>	<b>77.8</b>
<i>The category set is generated by NXTP based on the images.</i>						
210 (133)	<b>72.6</b>	<b>90.2</b>	<b>70.4</b>	84.6	41.5	75.7
794 (484)	72.5	90.4	70.3	<b>84.8</b>	<b>44.1</b>	<b>76.2</b>

Table 4. Ablation study on different category sets. The numbers in parentheses denote the number of foreground thing classes.

afforded by a broader category set.

To address the limitations of category set annotations, we generate category sets using the generative model NXTP [65] based on images and use DeepSeek-R1 [14] to classify foreground and background categories. As shown in Tab. 4, our model still achieves strong performance.

**Ablation study on different P-VLMs.** Tab. 5 shows that using ViT-B/16 as the retrieval model, our method performs well in mask classification, outperforming the P-VLM (ViT-B/16 with CLIPSelf). This improvement effectively mitigates foreground bias by enhancing the distinction between foreground and background regions. Using ViT-L/14 further improves performance across all metrics.

#### 4.4. Zero-shot comparisons at different data scales

To investigate the effect of data scale on VLM performance, we select three training sets of varying sizes from the SA-1B [26] dataset: 100K, 1.1M, and 5.5M seen sam-

P-VLM	Boxes		Masks-T		Masks-S	
	Top1	Top5	Top1	Top5	Top1	Top5
ViT-B/16	70.0	88.6	68.0	83.0	43.0	75.3
ViT-L/14	<b>73.4</b>	<b>90.5</b>	<b>71.0</b>	<b>84.8</b>	<b>45.6</b>	<b>77.8</b>

Table 5. Ablation study on different P-VLMs utilized for retrieving categories for unlabeled regions.

ples. We train our approach and the competing approach CLIPSelf on these subsets using  $8 \times A40$  GPUs and batch size of 48 per GPU. The resulting accuracy lines, presented in Fig. 6, reflect the performance of each model on the COCO and ADE20K benchmarks for open-vocabulary dense prediction tasks. As the size of the training set increases, the performance improvement of CLIPSelf slows down, while DenseVLM continues to show consistent improvement, demonstrating its promising scalability. The ability of our DenseVLM to rely on unlabeled data enables it to scale effectively to larger datasets.

#### 4.5. Extending to different backbones

To demonstrate the generality of our method, we extend it to several backbones, including ViT-B/16, R50x4, and ConvNeXt-B from OpenCLIP [4], as well as ViT-B/16 from EVA-CLIP [53]. From Tab. 6, we observe that CLIPSelf [57] is more effective with ViT-based architectures but is less effective with CNN-based architectures, particularly showing degraded performance with R50x4 compared to the baseline. In contrast, our approach yields consistent performance gains across all architectures, thereby broadening its applicability and potential. Furthermore, DenseVLM narrows the performance gap between powerful models like

Backbones	VLMs	Boxes	Masks-T	Masks-S
ViT-L/14	CLIPSelf	75.2	73.1	44.5
ViT-B/16	OpenCLIP	49.8	51.9	29.2
ViT-B/16*	CLIPSelf	67.6	64.4	44.5
ViT-B/16*	DenseVLM	<b>71.9</b>	<b>70.0</b>	<b>47.8</b>
ViT-B/16	EVA-CLIP	44.3	44.7	26.2
ViT-B/16	CLIPSelf	69.1	66.7	41.7
ViT-B/16	DenseVLM	<b>73.4</b>	<b>71.0</b>	<b>45.6</b>
R50x4	OpenCLIP	59.2	50.5	39.1
R50x4	CLIPSelf	59.1	49.9	37.2
R50x4	DenseVLM	<b>65.6</b>	<b>55.2</b>	<b>40.9</b>
ConvNeXt-B	OpenCLIP	57.5	48.0	31.1
ConvNeXt-B	CLIPSelf	62.6	57.6	41.6
ConvNeXt-B	DenseVLM	<b>67.1</b>	<b>63.4</b>	<b>43.5</b>

Table 6. Results and comparisons of different backbones. \* indicates the model initialized by OpenCLIP [4].

Method	A-150	A-847	PC-459
PACL [48]	31.4	-	-
OVSeg [36]	24.8	7.1	11.0
MAFT [21]	29.1	10.1	12.6
SED [60]	31.6	11.4	18.6
SCAN [40]	30.8	10.8	13.2
CAT-Seg+CLIPSelf [57]	29.7	10.1	-
CAT-Seg+FineCLIP [23]	32.4	12.2	-
SAN [62]	27.4	10.0	13.0
SAN+DenseVLM	<b>29.5</b> <sub>+2.1</sub>	<b>10.4</b> <sub>+0.4</sub>	<b>15.6</b> <sub>+2.6</sub>
CAT-Seg [5]	31.4	11.7	18.4
CAT-Seg+DenseVLM	<b>34.1</b> <sub>+2.7</sub>	<b>12.2</b> <sub>+0.5</sub>	<b>18.7</b> <sub>+0.3</sub>

Table 7. Results on open-vocabulary semantic segmentation.

ViT-L/14 fine-tuned by CLIPSelf and lightweight models, with some lightweight models even outperforming the more powerful ones on the Masks-S metric.

#### 4.6. Application to open-vocabulary dense tasks

**Experiment settings.** To evaluate the performance of the proposed DenseVLM in downstream tasks, we use DenseVLM as the backbone for open-vocabulary dense prediction tasks, including object detection and segmentation. To ensure fairness, our DenseVLM models are trained on the COCO train2017 split with an input resolution of  $512 \times 512$ .

**Open-vocabulary semantic segmentation.** We apply DenseVLM models initialized with OpenAI CLIP [50] to open-vocabulary semantic segmentation, including SAN [62] with a frozen backbone and Cat-Seg [5] with a fine-tuned backbone. The models are trained on COCO-Stuff [1] and evaluated on the ADE20K [71] (ADE-150 and ADE-847) and PASCAL Context [47] (PC-459) datasets, using the mean Intersection-over-Union (mIoU) metric. As shown in Tab. 7, DenseVLM consistently improves performance across all evaluation datasets, further enhancing the state-of-the-art performance.

(a) OV-COCO benchmark

Method	Backbone	$AP_{50}^{novel}$	$AP_{50}^{base}$	$AP_{50}$
VLDet [37]	RN50	32.0	50.6	45.8
F-VLM [27]	RN50	28.0	-	39.6
BARON-Cap [56]	RN50	33.1	54.8	49.1
CORA [58]	RN50	35.1	35.5	35.4
RO-ViT [25]	ViT-B/16	30.2	-	41.5
RO-ViT [25]	ViT-L/16	33.0	-	47.7
F-ViT+CLIPSelf [57]	ViT-B/16	25.4	40.9	36.8
F-ViT+FineCLIP [23]	ViT-B/16	29.8	45.9	41.7
F-ViT [57]	ViT-B/16	17.5	41.0	34.9
F-ViT+DenseVLM	ViT-B/16	<b>33.1</b> <sub>+15.6</sub>	<b>52.5</b> <sub>+11.5</sub>	<b>47.4</b> <sub>+12.5</sub>

(b) OV-LVIS benchmark

Method	Backbone	$mAP_r$	$mAP_c$	$mAP$
VLDet [37]	RN50	21.7	29.8	30.1
BARON-Cap [56]	RN50	22.6	27.6	27.6
F-VLM [27]	RN50	18.6	-	24.2
COR [58]	RN50x4	22.2	-	-
RO-ViT [25]	ViT-B/16	28.0	-	30.2
F-ViT+CLIPSelf [57]	ViT-B/16	10.6	7.6	9.3
F-ViT+FineCLIP [23]	ViT-B/16	10.4	8.0	9.5
F-ViT [57]	ViT-B/16	11.5	12.3	15.4
F-ViT+DenseVLM	ViT-B/16	<b>23.9</b> <sub>+12.4</sub>	<b>18.4</b> <sub>+6.1</sub>	<b>21.4</b> <sub>+6.0</sub>

Table 8. Results on open-vocabulary object detection.

**Open-vocabulary object detection.** Building on previous work [57], we construct open-vocabulary object detectors using the F-ViT architecture, a two-stage detector built on frozen ViTs from EVA-CLIP [53]. As shown in Tab. 8, we evaluate performance on the OV-COCO [3] benchmark by reporting box AP at IoU 0.5 for base, novel, and overall categories ( $AP_{50}^{novel}$ ,  $AP_{50}^{base}$ ,  $AP_{50}$ ), and on OV-LVIS [15] benchmark by mean AP for rare ( $mAP_r$ ), common ( $mAP_c$ ), and all categories ( $mAP$ ). Replacing the frozen EVA-CLIP ViTs with DenseVLM models leads to clear performance improvements on both benchmarks, achieving competitive performance relative to previous methods.

## 5. Conclusions

In this paper, we present DenseVLM, a framework designed to mitigate foreground bias in region-level vision-language alignment. DenseVLM can directly replace the original VLMs in open-vocabulary object detection and image segmentation methods, demonstrating consistently clear performance improvements over the baseline models. Furthermore, we validate DenseVLM’s promising scaling ability by exploring efficient region retrieval and decoupled alignment, successfully implementing DenseVLM to the training data from the SA-1B dataset. Overall, DenseVLM offers a generalizable solution for improving dense representations in vision-language models across various backbones, which is essential for open-vocabulary dense prediction tasks.

## Acknowledgments

This work was partially funded by NSFC (No. 62225604, 62276145) and the Science and Technology Support Program of Tianjin, China (No. 23JCZDJC01050). This work was also supported by the Supercomputing Center of Nankai University.

## References

- [1] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocomstuff: Thing and stuff classes in context. In *CVPR*, pages 1209–1218, 2018. 3, 8
- [2] Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 20(1):38–56, 2023. 3
- [3] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 2, 8
- [4] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *CVPR*, pages 2818–2829, 2023. 3, 5, 7, 8
- [5] Seokju Cho, Heeseong Shin, Sunghwan Hong, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. In *CVPR*, 2024. 1, 2, 3, 8
- [6] Ian Covert, Tony Sun, James Zou, and Tatsunori Hashimoto. Locality alignment improves vision-language models. *arXiv preprint arXiv:2410.11087*, 2024. 1, 3
- [7] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *CVPR*, pages 11583–11592, 2022. 2
- [8] Zheng Ding, Jieke Wang, and Zhuowen Tu. Open-vocabulary panoptic segmentation with maskclip. *arXiv preprint arXiv:2208.08984*, 2022. 2, 3
- [9] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [10] Alex Fang, Albin Madappally Jose, Amit Jain, Ludwig Schmidt, Alexander Toshev, and Vaishaal Shankar. Data filtering networks. *arXiv preprint arXiv:2309.17425*, 2023. 5
- [11] Zhiyu Fang, Xiaobin Zhu, Chun Yang, Hongyang Zhou, Jingyan Qin, and Xu-Cheng Yin. Aligning enhanced feature representation for generalized zero-shot learning. *Science China Information Sciences*, 68(2):122102, 2025. 1
- [12] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *ECCV*, pages 540–557, 2022. 1, 2
- [13] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *ICLR*, pages 1–20, 2022. 1, 2, 3
- [14] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 7
- [15] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 2, 8
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017. 3
- [17] Linlan Huang, Xusheng Cao, Haori Lu, and Xialei Liu. Class-incremental learning with clip: Adaptive representation adjustment and parameter fusion. In *ECCV*, pages 214–231. Springer, 2024. 1
- [18] Linlan Huang, Xusheng Cao, Haori Lu, Yifan Meng, Fei Yang, and Xialei Liu. Mind the gap: Preserving and compensating for the modality gap in clip-based continual learning, 2025. 1
- [19] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916, 2021. 3
- [20] Zihao Jia, Shengkun Sun, Guangcan Liu, and Bo Liu. Mssd: multi-scale self-distillation for object detection. *Visual Intelligence*, 2(1):8, 2024. 1
- [21] Siyu Jiao, Yunchao Wei, Yaowei Wang, Yao Zhao, and Humphrey Shi. Learning mask-aware clip representations for zero-shot segmentation. *NeurIPS*, 36:35631–35653, 2023. 3, 8
- [22] Siyu Jiao, Hongguang Zhu, Jiannan Huang, Yao Zhao, Yunchao Wei, and Humphrey Shi. Collaborative vision-text representation optimizing for open-vocabulary segmentation. *arXiv preprint arXiv:2408.00744*, 2024. 3
- [23] Dong Jing, Xiaolong He, Yutian Luo, Nanyi Fei, Wei Wei, Huiwen Zhao, Zhiwu Lu, et al. Fineclip: Self-distilled region-based clip for better fine-grained understanding. *NeurIPS*, 37:27896–27918, 2024. 5, 8
- [24] Cijo Jose, Théo Moutakanni, Dahyun Kang, Federico Baldassarre, Timothée Darcet, Hu Xu, Daniel Li, Marc Szafraniec, Michaël Ramamonjisoa, Maxime Oquab, et al. Dinov2 meets text: A unified framework for image-and pixel-level vision-language alignment. In *CVPR*, pages 24905–24916, 2025. 3
- [25] Dahun Kim, Anelia Angelova, and Weicheng Kuo. Region-aware pretraining for open-vocabulary object detection with vision transformers. In *CVPR*, pages 11144–11154, 2023. 1, 3, 8
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, pages 4015–4026, 2023. 2, 7
- [27] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. F-vlm: Open-vocabulary object detection upon frozen vision and language models. In *ICLR*, 2022. 1, 2, 8
- [28] Mengcheng Lan, Chaofeng Chen, Yiping Ke, Xinjiang Wang, Litong Feng, and Wayne Zhang. Clearclip: Decomposing clip representations for dense vision-language inference. In *ECCV*, 2024. 3

- [29] Mengcheng Lan, Chaofeng Chen, Yiping Ke, Xinjiang Wang, Litong Feng, and Wayne Zhang. Proxyclip: Proxy attention improves clip for open-vocabulary segmentation. In *ECCV*, 2024. 3
- [30] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989. 2
- [31] Jiaao Li, Yixiang Huang, Ming Wu, Bin Zhang, Xu Ji, and Chuang Zhang. Clip-sp: Vision-language model with adaptive prompting for scene parsing. *Computational Visual Media*, 10(4):741–752, 2024. 2
- [32] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *CVPR*, pages 10965–10975, 2022. 1, 3
- [33] Yunheng Li, ZhongYu Li, Quansheng Zeng, Qibin Hou, and Ming-Ming Cheng. Cascade-clip: Cascaded vision-language embeddings alignment for zero-shot semantic segmentation. In *ICML*, pages 28243–28258, 2024. 1, 2
- [34] Zheng Li, Xiang Li, Xinyi Fu, Xin Zhang, Weiqiang Wang, Shuo Chen, and Jian Yang. Promptkd: Unsupervised prompt distillation for vision-language models. In *CVPR*, pages 26617–26626, 2024. 1
- [35] Zheng Li, Yibing Song, Ming-Ming Cheng, Xiang Li, and Jian Yang. Advancing textual prompt learning with anchored attributes. *arXiv preprint arXiv:2412.09442*, 2024. 1
- [36] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*, pages 7061–7070, 2023. 1, 2, 8
- [37] Chuang Lin, Peize Sun, Yi Jiang, Ping Luo, Lizhen Qu, Gholamreza Haffari, Zehuan Yuan, and Jianfei Cai. Learning object-language alignments for open-vocabulary object detection. In *ICLR*, 2032. 8
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [39] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 1, 3
- [40] Yong Liu, Sule Bai, Guanbin Li, Yitong Wang, and Yansong Tang. Open-vocabulary segmentation with semantic-assisted calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3491–3500, 2024. 8
- [41] Yong Liu, Sule Bai, Guanbin Li, Yitong Wang, and Yansong Tang. Open-vocabulary segmentation with semantic-assisted calibration. In *CVPR*, 2024. 2
- [42] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 5
- [43] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [44] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, pages 13–23, 2019. 3
- [45] Hao Ma, Ming Li, Jingyuan Yang, Or Patashnik, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Clip-flow: Decoding images encoded in clip space. *Computational Visual Media*, 10(6):1157–1168, 2024. 3
- [46] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *ECCV*, pages 728–755, 2022. 1, 2
- [47] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, pages 891–898, 2014. 8
- [48] Jishnu Mukhoti, Tsung-Yu Lin, Omid Poursaeed, Rui Wang, Ashish Shah, Philip HS Torr, and Ser-Nam Lim. Open vocabulary semantic segmentation with patch aligned contrastive learning. In *CVPR*, pages 19413–19423, 2023. 3, 8
- [49] Muhammad Ferjad Naeem, Yongqin Xian, Xiaohua Zhai, Lukas Hoyer, Luc Van Gool, and Federico Tombari. Silc: Improving vision language pretraining with self-distillation. In *ECCV*, pages 38–55. Springer, 2024. 3
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 1, 2, 3, 8
- [51] Kanchana Ranasinghe, Brandon McKinzie, Sachin Ravi, Yinfei Yang, Alexander Toshev, and Jonathon Shlens. Perceptual grouping in contrastive vision-language models. In *ICCV*, pages 5571–5584, 2023. 3
- [52] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, pages 25278–25294, 2022. 3
- [53] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023. 1, 2, 3, 5, 6, 7, 8
- [54] Feng Wang, Jieru Mei, and Alan Yuille. Sclip: Rethinking self-attention for dense vision-language inference. In *ECCV*, pages 315–332, 2025. 3
- [55] Xiao Wang, Guangyao Chen, Guangwu Qian, Pengcheng Gao, Xiao-Yong Wei, Yaowei Wang, Yonghong Tian, and Wen Gao. Large-scale multi-modal pre-trained models: A comprehensive survey. *Machine Intelligence Research*, 20(4):447–482, 2023. 3
- [56] Size Wu, Wenwei Zhang, Sheng Jin, Wentao Liu, and Chen Change Loy. Aligning bag of regions for open-vocabulary object detection. In *CVPR*, pages 15254–15264, 2023. 2, 8

- [57] Size Wu, Wenwei Zhang, Lumin Xu, Sheng Jin, Xiangtai Li, Wentao Liu, and Chen Change Loy. Clipself: Vision transformer distills itself for open-vocabulary dense prediction. *ICLR*, 2024. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [58] Xiaoshi Wu, Feng Zhu, Rui Zhao, and Hongsheng Li. Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In *CVPR*, pages 7031–7040, 2023. [2](#), [8](#)
- [59] Monika Wysoczańska, Oriane Siméoni, Michaël Ramamonjisoa, Andrei Bursuc, Tomasz Trzciniński, and Patrick Pérez. Clip-dinoiser: Teaching clip a few dino tricks for open-vocabulary semantic segmentation. In *ECCV*, pages 320–337. Springer, 2024. [3](#)
- [60] Bin Xie, Jiale Cao, Jin Xie, Fahad Shahbaz Khan, and Yanwei Pang. Sed: A simple encoder-decoder for open-vocabulary semantic segmentation. In *CVPR*, 2024. [3](#), [8](#)
- [61] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *CVPR*, pages 2955–2966, 2023. [1](#), [2](#)
- [62] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *CVPR*, pages 2945–2954, 2023. [1](#), [2](#), [3](#), [8](#)
- [63] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. San: Side adapter network for open-vocabulary semantic segmentation. *IEEE TPAMI*, pages 1–17, 2023. [2](#)
- [64] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. *NeurIPS*, 2023. [2](#)
- [65] Kaiyu Yue, Bor-Chun Chen, Jonas Geiping, Hengduo Li, Tom Goldstein, and Ser-Nam Lim. Object recognition as next token prediction. In *CVPR*, pages 16645–16656, 2024. [2](#), [3](#), [7](#)
- [66] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *CVPR*, pages 14393–14402, 2021. [1](#), [2](#)
- [67] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, pages 11975–11986, 2023. [5](#)
- [68] Haotian Zhang, Pengchuan Zhang, Xiaowei Hu, Yen-Chun Chen, Liunian Li, Xiyang Dai, Lijuan Wang, Lu Yuan, Jenq-Neng Hwang, and Jianfeng Gao. Glipv2: Unifying localization and vision-language understanding. *NeurIPS*, 35: 36067–36080, 2022. [1](#), [3](#)
- [69] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *CVPR*, pages 16793–16803, 2022. [1](#), [2](#), [3](#), [5](#)
- [70] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. [5](#)
- [71] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 127: 302–321, 2019. [3](#), [8](#)
- [72] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *ECCV*, page 696–712, 2022. [1](#), [3](#)
- [73] Yuanbing Zhu, Bingke Zhu, Zhen Chen, Huan Xu, Ming Tang, and Jinqiao Wang. Mrowseg: Breaking the resolution curse of vision-language models in open-vocabulary semantic segmentation. *arXiv preprint arXiv:2408.14776*, 2024. [2](#)