

Toward Modality Gap: Vision Prototype Learning for Weakly-supervised Semantic Segmentation with CLIP

Zhongxing Xu^{1*}, Feilong Tang^{1*}, Zhe Chen², Yingxue Su²,
Zhiyi Zhao³, Ge Zhang², Jionglong Su², Zongyuan Ge^{1†}

¹AIM Lab, Faculty of IT, Monash University

²Xi'an Jiaotong-Liverpool University

³University of Edinburgh

Zhongxing.Xu@monash.edu

Abstract

The application of Contrastive Language-Image Pre-training (CLIP) in Weakly Supervised Semantic Segmentation (WSSS) research powerful cross-modal semantic understanding capabilities. Existing methods attempt to optimize input text prompts for improved alignment of images and text, by finely adjusting text prototypes to facilitate semantic matching. Nevertheless, given the modality gap between text and vision spaces, the text prototypes employed by these methods have not effectively established a close correspondence with pixel-level vision features. In this work, our theoretical analysis indicates that the inherent modality gap results in misalignment of text and region features, and that this gap cannot be sufficiently reduced by minimizing contrast loss in CLIP. To mitigate the impact of the modality gap, we propose a Vision Prototype Learning (VPL) framework, by introducing more representative vision prototypes. The core of this framework is to learn class-specific vision prototypes in vision space with the help of text prototypes, for capturing high-quality localization maps. Moreover, we propose a regional semantic contrast module that contrasts regions embedding with corresponding prototypes, leading to more comprehensive and robust feature learning. Experimental results show that our proposed framework achieves state-of-the-art performance on two benchmark datasets.

Introduction

Semantic segmentation serves as a fundamental task in the field of computer vision (Tang et al. 2023; Wang et al. 2022; Xu et al. 2024b; Hu et al. 2025, 2024b; Yang et al. 2023; Xiong et al. 2024; Trinh et al. 2024; Tang et al. 2024a; Hu et al. 2024a; Zhang et al. 2024b; Wan et al. 2024). Weakly Supervised Semantic Segmentation (WSSS) has gained popularity in the research community. It learns from weak labels such as image-level labels (Lee, Kim, and Yoon 2021), scribbles (Lin et al. 2016), or bounding boxes (Lee et al. 2021), instead of pixel-level annotations. Most WSSS approaches utilize CAM (Zhou et al. 2016) to provide localization cues for target objects, thereby mapping visual con-

*These authors contributed equally.

†Corresponding author: Zongyuan Ge

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

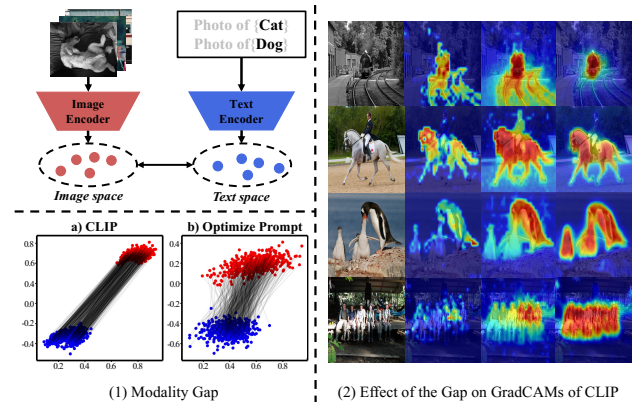


Figure 1: The main idea proposed in this paper is to reduce the impact of the modality gap by learning vision prototypes. We show the modality gap between paired image embeddings and text prototypes. Even though the Optimized prompt minimizes contrastive loss between prototypes and object region, the modality gap is still not sufficiently reduced (as shown in (a)), which results in the text prototypes failing to accurately capture the relevant region of the target object. In contrast, we propose VPL, which ensures accurate activation of the complete object region (b).

cepts to pixel regions. After training with 400 million image-text pairs, Contrastive Language-Image Pre-training (Radford et al. 2021) has established a strong association between vision concepts and text descriptions, demonstrating excellent cross-modal semantic matching capabilities. Inspired by transfer learning, several WSSS methods (Yang and Gong 2024) apply powerful text prototypes encoded by the CLIP model to generate high-quality CAMs.

The key to WSSS is to generate CAM that effectively covers the complete object. Recent weakly supervised research (Lin et al. 2023; Xu et al. 2023; Liu et al. 2023; Yang and Gong 2024) primarily focus on constructing more adaptive text prototypes by optimizing text cues, with the aim to establish stronger correlations with semantic object regions for dense localization. Customized text prompting strategies can enhance the semantic alignment between object regions and

their corresponding text prototypes to a certain extent, but a significant semantic gap still persists. This gap is due to the inherent cone effect of CLIP, which causes the modality gap between text and vision spaces (Liang et al. 2022). As shown in Fig 1 (a), even after minimizing the contrastive loss between the text prototype and the semantic object region, the cone effect remains significant.

The semantic misalignment caused by the modality gap affects the discriminative ability of the text prototypes in two aspects. On the one hand, our theoretical analysis indicates that text prototypes contain inactive and redundant information. Since CLIP was trained with millions of captions and concepts, its text embeddings contain a large number of features that may not be relevant to a specific target domain (Hu et al. 2024c). However, the categories in the target domain usually require the activation of only a limited set of specific category features, resulting in a large amount of unnecessary text information in the text prototypes. On the other hand, text prototypes lack vision details such as lighting, color, and texture. This absence weakens their ability to match the detailed visual features, leading to a deficiency category representativeness of text prototypes. As shown in Fig 1 (b), although the optimized text prototypes show effectiveness, they still struggle to capture more complete and accurate target object areas. Moreover, they erroneously activate similar or highly co-occurring categories (e.g., `railroad` and `train` in Fig. 1 (b)). Therefore, the text prototypes are affected by the modality gap leading to suboptimal segmentation performance of the model.

To address these challenges, we propose a **Vision Prototype Learning (VPL)** framework to reduce the impact of the modality gap by introducing more representative vision prototypes. We theoretically demonstrate that the optimal category prototypes for vision tasks should originate from the vision space. Moreover, our analysis reveals that the modality gap between text and vision spaces obtained by CLIP is inherent, confirming that optimal vision prototypes cannot be achieved in text space. Our VPL framework is comprised of two main phases. In the first phase, instead of taking the conventional approach of directly using text prototypes, we use text prototypes with KL-divergence as a constraint to precisely learn vision prototypes in the vision space through gradient descent. These vision prototypes are then utilized to regenerate more accurate pseudo-masks, as shown in Fig. 1 (b). In the second phase, we use the refined pseudo-masks to effectively supervise the decoder, transforming high-level feature information into accurate pixel-level predictions. However, imperfect mask labels provide ambiguous knowledge to the segmentation network. To address the issues, we propose regional semantic contrast that contrasts corresponding mask region embedding and vision prototypes, driving the network to learn more comprehensive and robust object embeddings and enhancing adaptability to the modality gap.

In the PASCAL VOC 2012 (Everingham et al. 2010) and MS COCO 2014 (Lin et al. 2014) datasets, we evaluate our method in various WSSS settings, where our approach achieves state-of-the-art performance. The key contributions of our research can be summarized as follows:

- We are the first to investigate the impact of modality gap in weakly-supervised semantic segmentation with CLIP and introduce the Vision Prototype Learning framework to alleviate this issue.
- We theoretically demonstrate the inherent modality gap in CLIP and propose to obtain vision prototypes from the vision space with help of the text prototypes.
- We introduce a regional semantic contrast module to enhance the alignment between object regions and vision prototypes. Experiments demonstrate the efficacy of our framework with state-of-the-art performance on mainstream benchmarks.

Related Work

Weakly Supervised Semantic Segmentation as an important topic in computer vision (Ye et al. 2024; Chen et al. 2024; Zhang et al. 2024d; Wei and Zhang 2024; Yuan et al. 2024; Zhang et al. 2024c), can generally be divided into single-stage and multi-stage learning processes. Single-stage methods (Wang et al. 2023; Wu et al. 2024a) use image-level labels to train the segmentation network in an end-to-end manner. In contrast, multi-stage methods (Rong et al. 2023; Wu et al. 2023; Zhou et al. 2022; Cho et al. 2024) first generate segmentation seeds and then use them as pseudo-labels to train off-the-shelf segmentation network for better performance. In order to generate high-quality seeds, recent research has proposed solutions such as adversarial erasure (Yoon et al. 2022), region growing (Peng et al. 2023), localizing attention map (Xu et al. 2024a), and exploring boundary constraints (Ru et al. 2023).

Contrastive Language-Image Pre-training achieves strong generalization ability through training with a large number of vision-text pairs. Although CLIP excels in cross-modal learning, (Liang et al. 2022) indicates that there is an inherent gap between different data modalities within the CLIP model, leading to insufficient alignment between text and visual representations. Modality gap have been intensively studied in fields such as few-shot learning (Ouali et al. 2023), zero-shot learning (Qian, Xu, and Hu 2023), and domain adaptation (Hu et al. 2024c). To mitigate this issue, SuS-X (Udandarao, Gupta, and Albanie 2023) introduces additional images to provide richer visual information. In recent developments, CLIP has been widely applied in WSSS. CLIMS (Xie et al. 2022) introduces a loss function derived from CLIP to guide a network to produce high-quality CAM. CLIP-ES (Lin et al. 2023) has developed GradCAM and Class-Aware Attention Affinity (CAA) techniques for directly creating CAM using CLIP. Meanwhile, WeCLIP (Zhang et al. 2024a) leverages a frozen CLIP model to extract semantic features and combines it with a new decoder and refinement module (RFM) to generate and correct pseudo labels. In addition to these approaches, FMA (Yang and Gong 2024) proposes two sets of task-specific learnable text prompts to capture category semantic knowledge relevant to segmentation, thereby constructing more representative text category prototypes.

Preliminary

Given each training image $x_i \in \mathbb{R}^{w \times h \times 3}$ in the dataset $\{x_i\}_{i=1}^I$, where w and h are the width and height of the images, respectively, and I is the total number of images. Each image is associated with only an image-level label $\{z_n\}_{n=1}^N$ in which N is the pre-specified categories, and their corresponding vision and text representations can be extracted as:

$$\mathbf{x}_i = E_v(x_i); \quad \mathbf{z}_n = E_t(z_n); \quad M_n = \text{GradCAM}(\mathbf{x}_i, \mathbf{x}_n), \quad (1)$$

where $E_v(\cdot)$ and $E_t(\cdot)$ denote the image and text encoder in the pre-trained CLIP model, respectively. $\mathbf{x}_i \in \mathbb{R}^{S \times H \times D}$ and $\mathbf{z}_n \in \mathbb{R}^D$ represent the vision features and the text features with D channels. \mathbf{x}_i has a spatial size of $S \times H$, and M_n denotes a GradCAM. See Appendix A for more details.

Modality Gap in CLIP. Given a batch of B (v_i, t_i) pairs, CLIP learns to maximize the cosine similarity of the image and text embeddings of the B real pairs in the batch while minimizing the cosine similarity of the embeddings of the $B^2 - B$ incorrect pairs. The model learns two encoders by minimizing the contrastive loss ℓ , which can be written as:

$$\ell(\mathbf{v}_i, \mathbf{t}_i) = \sum_i -\log \frac{\exp(\mathbf{v}_i^\top \mathbf{t}_i / \tau)}{\sum_j \exp(\mathbf{v}_i^\top \mathbf{t}_j / \tau)} - \log \frac{\exp(\mathbf{t}_i^\top \mathbf{v}_i / \tau)}{\sum_j \exp(\mathbf{t}_i^\top \mathbf{v}_j / \tau)}. \quad (2)$$

where \mathbf{v}_i and \mathbf{t}_i represent the L2-normalized embeddings of the image and text in the i -th pair, respectively. τ is a learned temperature parameter to scale the logits. The cross-entropy loss compares the similarity between image and text embeddings, focusing on cross-modal association and ranking.

A significant difference between inter-modal and intra-modal distributions in CLIP has been observed (Liang et al. 2022; Udandarao, Gupta, and Albanie 2023). The Euclidean distance between modalities, *i.e.*, $\sum_i \|\mathbf{v}_i - \mathbf{t}_i\|_2^2$, is related to the magnitude of temperature. However, the small τ does not pull the text and vision space together (the proof in Appendix B.1). The CLIP model focuses on high-level associations between images and texts rather than fine-grained vision features, and it does not merge images and texts into a unified, indistinguishable representation space, as shown in Fig. 1 (b). This inability to map to the visual features required for precise segmentation means that the prototypes from the text space may not capture the nuances of the vision space, resulting in degraded performance for segmentation tasks.

Methodology

Learning Vision Prototypes from Text Supervision

To mitigate the impact of the modality gap, we propose to introduce more representative prototypes to better capture the spatial coverage of semantic object regions. We utilize mask labels $\{m_i\}$ for the standard semantic segmentation supervised learning in vision space as follows:

$$\ell(\mathbf{x}_i, m_{i,k}) = \min_W \sum_{i,k} -\log \left(\frac{\exp(\mathbf{x}_{i,k}^\top \mathbf{w}_{m_{i,k}} / \tau_I)}{\sum_n \exp(\mathbf{x}_{i,k}^\top \mathbf{w}_n / \tau_I)} \right), \quad (3)$$

where \mathbf{w}_n denotes the learnable vision classifier, and $W = \{\mathbf{w}_n\}_{n=1}^N \in \mathbb{R}^{d \times N}$, $\mathbf{x}_{i,k}$ denotes the k -th pixel representation in image i . τ_I is the temperature parameter for the optimization with vision data. Minimizing the objective function in Eq. (3) requires the adjustment of \mathbf{w}_n . This causes the **decision boundary** of \mathbf{w}_n to generate higher scores in the feature space $\{\mathbf{x}_i\}$, with the aim to maximally reflect the semantic information and category distribution of the images.

We raise a question: Is it possible to find the optimal vision prototypes within the text space?

Hypothesis 1. We decompose the class prototype as:

$$\mathbf{z}_n = \sqrt{\epsilon} \mathbf{z}_n^x + \sqrt{1 - \epsilon} \mathbf{z}_n^\perp, \quad (4)$$

$$\text{s.t.} \quad \|\mathbf{z}_n^x\| = \|\mathbf{z}_n^\perp\| = 1, \quad \mathbf{z}_n^{x\top} \mathbf{z}_n^\perp = 0, \quad 0 \leq \epsilon \leq 1,$$

where \mathbf{z}_n^x is from the vision space spanned by $\{\mathbf{x}_i\}$, and \mathbf{z}_n^\perp represents the component from the orthogonal subspace with the condition that $\mathbf{z}_n^{x\top} \mathbf{z}_n^\perp = 0$, and they have unit norms. We note that $\lim_{\epsilon \rightarrow 1} \mathbf{z}_n = \mathbf{z}_n^x$. In this case, \mathbf{z}_n retains more vision information.

Remark: The magnitude of ϵ determines the overlap between two components. Our hypothesis suggests that with this decomposition, if the vision space is covered by the text space, such that the class prototype \mathbf{z}_n contains all information from both modalities, then it is possible to recover optimal segmentation performance. The details are as follows.

Proposition 1. Define $p_{i,k}^n$ and $p_{i,k}^{n'}$ respectively denote the predicted probabilities for each pixel k obtained by the vision prototypes W^* and text prototypes Z . Then, we have:

$$p_{i,k}^{n'} = \frac{\exp(\mathbf{x}_{i,k}^\top \mathbf{z}_n / \tau_T)}{\sum_n \exp(\mathbf{x}_{i,k}^\top \mathbf{z}_n / \tau_T)} = \frac{\exp(\sqrt{\epsilon} \mathbf{x}_{i,k}^\top \mathbf{w}_n^* / \tau_T)}{\sum_n \exp(\sqrt{\epsilon} \mathbf{x}_{i,k}^\top \mathbf{w}_n^* / \tau_T)}, \quad (5)$$

where τ_T represents the temperature in CLIP, while τ_I is the temperature for learning vision prototypes. If $\mathbf{z}_n^x = \mathbf{w}_n^*$, then $p_{i,k}^{n'} = p_{i,k}^n$ and $\tau_T = \sqrt{\epsilon} \tau_I$. See proof in Appendix B.3. \square

Remark: Proposition 1 suggests that if the text space encompasses the entire vision space, a larger temperature $\tau_T = \sqrt{\epsilon} \tau_I$ can help identify the optimal vision prototypes within the text space by smoothing the probability distribution and narrowing the modality gap.

Despite significant progress in integrating natural language and vision information, the text and vision spaces learned by CLIP are distinct, with a clear margin, as illustrated in Fig. 1 (a). It is difficult for the text space to cover the vision space. The assumptions (Hypothesis 1 and Proposition 1) suggesting that the text space can recover the optimal vision prediction are difficult to hold. This is due to the lack of sufficient detail in textual data to precisely describe the information in visual data (*e.g.*, colors and textures). Thus, the capability of text prototypes to provide precise guidance for pixel-level classification is limited. The lower bound for the modality gap is quantified in the following theorem.

Theorem 1 (The lower-bound of modality gap). *The image and text embeddings of CLIP are always on the unit sphere, with the Euclidean distance $\|Z - W^*\|^2$, being defined as:*

$$\Delta_{gap} = \|Z - W^*\|_F^2 \geq \mathbf{c}_\perp + \epsilon, \quad (6)$$

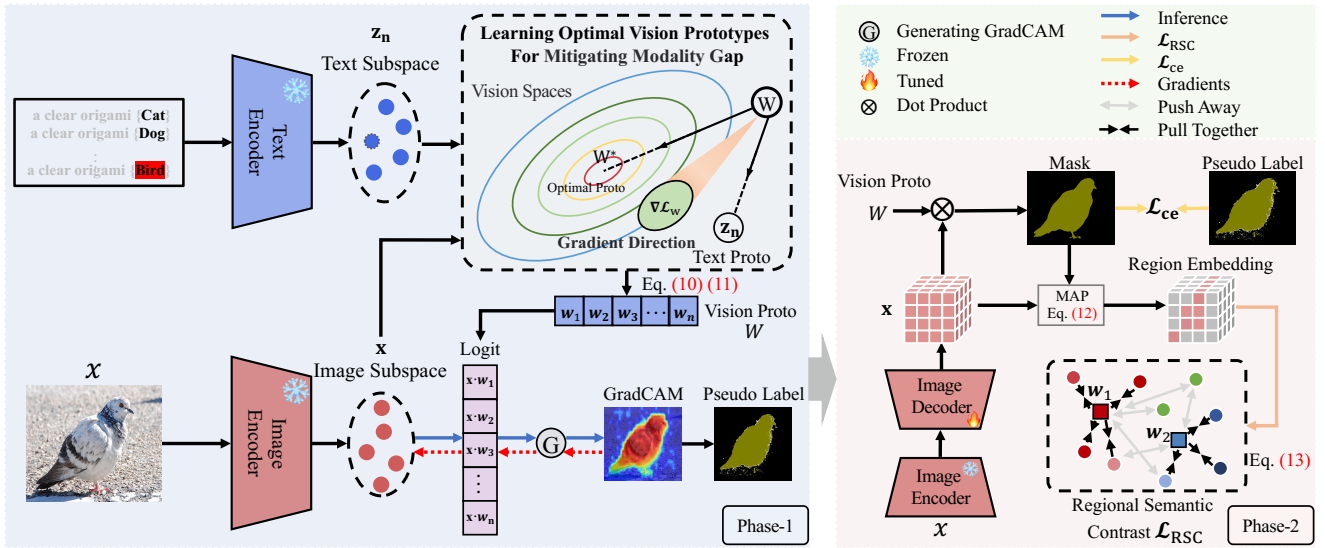


Figure 2: Overview of the proposed Weakly-supervised Vision Prototype Learning (VPL), which consists of two main components: (1) Learning the vision prototype W and (2) Regional semantic contrast (RSC). In phase (1), vision prototypes can be efficiently obtained by solving a convex optimization problem using gradient descent in Eq. (11). This ensures that it can align with vision data better to mitigate the impacts of the modality gap. The vision prototypes will then replace text prototypes to locate the target object and generate initial GradCAMs. In phase (2), we refine these GradCAMs to form pseudo-labels for supervising the decoder while CLIP encoders are frozen. Then, we obtain the masks from the vision prototypes and align them with specific region embeddings.

where c_{\perp} denote a constant vector, which is orthogonal to the embedding span of Z and W^* . And ϵ representing the alignment noise. See proof details in Appendix B.4. \square

Remark: Theorem 1 emphasizes the modality gap between the text prototype and the optimal vision prototype and consists of two components. c_{\perp} represents the distance to the irrelevant text space, i.e., redundant text dimensions. This implies that within the text space, there exist dimensions or information that are irrelevant to the vision modality. ϵ represents the approximation loss caused by the low-rank overlap between the text and vision spaces, which indicates that there may be a certain degree of overlap between the text and vision spaces. They share some information but cannot match each other completely, indicating the presence of partial semantic information missing in the text prototype. Thus, minimizing the distance between Z and W is difficult due to the inherent modality gap, which confirms our empirical observation.

To accurately capture vision semantic features, we obtain vision prototypes in the following Theorem 3. Since without precise pixel-level annotations, object masks in Eq. (3) can be generated by text prototypes using GradCAM. Our objective function is constructed by measuring the KL-divergence between the distributions of text prototypes and vision prototypes as:

$$\min_W \mathcal{L}(\mathcal{P}', W) = \sum_i \text{KL}(\mathcal{P}_i^{n'} \parallel \mathcal{P}_i^n), \quad (7)$$

where $\mathcal{P}_i^{n'}$ and \mathcal{P}_i^n respectively represent the distributions estimated by the text prototypes Z and the learnable vi-

sion prototypes W . This constraint aligns the model within a common representation space, reducing the modality gap.

Furthermore, we optimize the KL divergence in two ways: one is class-level regularization, where given anchor class n , the distribution across all categories can be computed. The other is instance-level distribution regularization, where the anchor class n is replaced with anchor \mathbf{x}_i .

Theorem 2. Assuming the norm of prototypes is bounded by η , i.e., $\forall n, \|\mathbf{w}_n\|_2 \leq \eta$, the distribution defined by the anchor class is an approximation of the distribution defined by the anchor example as:

$$\forall n, \frac{1}{\exp(2\kappa)} \mathcal{P}_{m_i, k} \leq \mathcal{P}_{i, j} \leq \exp(2\kappa) \mathcal{P}_{m_i, k}, \quad (8)$$

where $\kappa = 2\eta \|\mathbf{x}_i - \mathbf{w}_{m_i, k}\|_2$. When the intra-class distribution is compact, i.e., $\|\mathbf{x}_i - \mathbf{w}_{m_i, k}\|_2 \rightarrow 0$, this approximation becomes tight. See proof details in Appendix B.5. \square

Remark: Theorem 2 demonstrates that instance-level regularization, as opposed to class-level regularization, can help better capture variations in real data.

Theorem 3 (Learning Vision Prototypes). Assume $\mathcal{L}(\mathcal{P}', W)$ is a μ -strongly convex function in W , we have:

$$W^{t*} = \arg \min_W \mathcal{L}(\mathcal{P}', W); \quad W^* = \arg \min_W \mathcal{L}(m, W), \quad (9)$$

where m_i is the GradCAM as the pseudo mask distribution for \mathbf{x}_i , and is generated through text prototypes Z . We compute the Eq. (7) as $-\sum_{i, n, k} \mathcal{P}_{i, k}^{n'} \log(\mathcal{P}_{i, k}^n)$, and we compute the gradient direction to adjust W^{t*} by the standard gradi-

Algorithm 1: Learning the Optimal Vision Prototypes

Require: Input image set $\{x_i\}$, class names $\{z_n\}$, paired image text encoder E_v, E_t , vision prototype W^{I*} , loss function \mathcal{L} , iterations T_w , temperature τ_T, τ_I , learning rate α

- 1: $\mathbf{x}_i \leftarrow E_v(x_i), \mathbf{z}_n \leftarrow E_t(z_n)$ ▷ See Eq. (1)
 - 2: $W^{I*} \leftarrow \min_W \mathcal{L}(P^I, W)$ ▷ See Eq. (9)
 - 3: Calculate gradient $\text{Grad} = \nabla_{W^{I*}} \mathcal{L}(W)$ ▷ See Eq. (10)
 - 4: $W^{I*} \leftarrow W^{I*} - \alpha \nabla_{W^{I*}} \mathcal{L}(W)$ ▷ See Eq. (11)
 - 5: Calculate GradCAMs M_n using W^{I*} ▷ GradCAM calculation (See Eq. (1))
 - 6: **return** W^{I*}, M_n
-

ent descent. Subsequently, we have:

$$\nabla_{W^{I*}} \mathcal{L}(W) \leq \frac{2}{\mu} \langle P^I - Y, \log(\mathcal{P}_{W^{I*}}) - \log(\mathcal{P}_{W^*}) \rangle, \quad (10)$$

where $\nabla_{W^{I*}} \mathcal{L}(W)$ denote $\|W^{I*} - W^*\|^2$. Then, we update the W^{I*} to minimize the loss function:

$$W^{I*} \leftarrow W^{I*} - \alpha \nabla_{W^{I*}} \mathcal{L}(W) \quad (11)$$

where α is the learning rate, we iterate T_w times to gradually optimize the parameters. Update W^{I*} in the negative gradient direction to obtain a local optimal solution for the vision prototype. See proof details in Appendix B.6. \square

Remark: The core of our method is using KL-divergence and strong convexity to efficiently approach the optimal solution by gradient descent. By Proposition 1, the temperature τ_I during the learning of vision prototypes should be larger than τ_T in CLIP, and has been confirmed in our ablation study. Our experiments show that our vision prototype significantly improves the accuracy of target localization compared to the use of text prototypes with single prompt or ensemble prompts. As illustrated in Fig. 1 (b), the GradCAM generated by the vision prototypes focuses more on the target object rather than the background or irrelevant regions. The detailed algorithmic is given in Alg. 1.

Regional Semantic Contrast Learning. To learn the alignment between pixel-level features and prototypes, we supervise the mask decoder using post-processing GradCAM. However, relying solely on this loss is insufficient to fully exploit object regions as the noise in pseudo labels leads to errors in segmentation. Therefore, we propose a regional semantic contrast module to effectively contrast categorical object regions and vision prototypes by suppressing negative masks obtained from unrelated texts (*i.e.*, texts of negative pairs). Specifically, for the pixel-level dense embeddings $\mathbf{V} \in \mathbb{R}^{B \times D \times H \times S}$ from the decoder and the vision prototypes $W \in \mathbb{R}^{D \times N}$, we compute the feature-level localized image embeddings $P_{i,n} \in \mathbb{R}^D$ by masked average pooling (MAP) as follow:

$$P_{i,n} = \frac{\sum_{h,w} \mathbf{A}_{i,n,h,w} \cdot \mathbf{V}_{i,:,h,w}}{\sum_{h,w} \mathbf{A}_{i,n,h,w}}, \quad (12)$$

where B denotes a batch size. $\mathbf{A}_{i,n} = W_n^\top \mathbf{V}_i$ calculates the masks in the batch. Subsequently, we calculate the cosine similarity between vision prototypes and feature-level

localized image embeddings as $S_{i,n} = P_{i,n}^\top W_n$. RSC Loss is defined as follows:

$$\mathcal{L}_{\text{RSC}} = -\frac{1}{2B} \sum_i \log \frac{\exp(S_{i,i}/\tau)}{\sum_j \exp(S_{i,j}/\tau)} - \frac{1}{2B} \sum_i \log \frac{\exp(S_{i,i}/\tau)}{\sum_j \exp(S_{j,i}/\tau)}. \quad (13)$$

Final Loss. Our final loss function is defined by:

$$\mathcal{L} = \lambda_{\text{ce}} \mathcal{L}_{\text{ce}}(M, \mathbf{A}) + \lambda_{\text{RSC}} \mathcal{L}_{\text{RSC}}, \quad (14)$$

where \mathcal{L}_{ce} is the cross-entropy loss and \mathcal{L}_{RSC} is the region-level RSC loss. M denotes the mask obtained by post-processing the GradCAM. λ_{ce} and λ_{RSC} are coefficients.

Experiments

Datasets and Implementation Details

Datasets and Evaluation Metrics. Experiments are conducted on two benchmarks: PASCAL VOC 2012 (Everingham et al. 2010) with 21 classes and MS COCO 2014 (Lin et al. 2014) with 81 classes. For PASCAL VOC 2012, following (Wang et al. 2020; Lee, Kim, and Yoon 2021; Li et al. 2022), we use the augmented SBD (Hariharan et al. 2011) with 10,582 annotated images. We evaluate VPL in terms of i) the quality of generated pseudo masks on VOC 2012 train, and ii) semantic segmentation on VOC 2012 val/test and COCO 2014 val. Mean intersection over union (mIoU) (Long, Shelhamer, and Darrell 2015) is used as the metric in both cases. VOC 2012 test scores are obtained from the official evaluation server.

Implementation Details. In our experiments, we adopt the CLIP pre-trained model ViT-B-16 (Radford et al. 2021). The feature map used to generate CAM is the one before the last self-attention layer in ViT. We replace the class token with the average of remaining tokens to compute final logits. For the VOC 2012 dataset, phase 1 of our framework is learned by standard projected gradient descent, where the learning rate is 10 and the number of iterations is 3,000. τ_T is 0.01 in CLIP and τ_I for learning the vision prototypes is set to 0.03 in Section . CAMs are refined into pseudo masks using denseCRF (Krähenbühl and Koltun 2011). In phase 2, the batch size is set as 6 and the maximum iteration is set as 30,000. SGD optimizer is adopted to train only our decoder with a momentum of 0.9 and a weight decay of 1e-4. All experiments are conducted on 8 Nvidia 4090 GPUs. The loss coefficients λ_{ce} and λ_{RSC} are both set as 1 in Eq. 14.

Comparisons With State-of-the-Art Methods

Improved Localization Maps and Segmentation Results:

The proposed CLIP-CPAL does not modify the architecture of the CLIP network. It simply integrates VPL into the CLIP-ES and CLIP-CPAL methods. The results on VOC 2012 show improved localization maps, as shown in Table 2. For instance, incorporating VPL into CLIP-ES improves seed performance by 5.5% and pseudo mask performance by 2.5%. When VPL is added to the CLIP-CPAL model, there is a 5.9% gain in seed performance. Table 2 presents the performance of the semantic segmentation model trained

Method	Seed	Pseudo-Mask
IRN (Ahn, Cho, and Kwak 2019)	48.8	66.5
ESOL (Li et al. 2022)	53.6	68.7
MCTformer (Xu et al. 2022)	61.7	69.1
LPCAM (Chen and Sun 2023)	65.3	72.7
ACR (Kweon, Yoon, and Yoon 2023)	60.3	72.3
D2CAM (Wang et al. 2023)	58.0	71.4
Mat-Label (Wang et al. 2023)	62.3	72.9
FPR (Chen et al. 2023)	67.7	72.8
USAGE (Peng et al. 2023)	71.9	72.8
SFC (Zhao et al. 2024a)	64.7	73.7
SeCo (Yang et al. 2024)	74.8	76.5
SFC (Zhao et al. 2024b)	64.7	73.7
DuPL (Wu et al. 2024b)	73.5	75.1
PSDPM (Zhao et al. 2024b)	68.5	77.3
CLIP-ES (Lin et al. 2023)	70.8	75.0
+VPL (Ours)	76.3^{↑5.5}	78.4^{↑3.4}
CLIP-CPAL (Tang et al. 2024b)	71.9	75.8
+VPL (Ours)	77.8^{↑5.9}	80.1^{↑4.3}

Table 1: Comparisons between our method and the other WSSS methods. Evaluate mIoU (%) on the PASCAL VOC 2012 *train* set at levels: CAM and Pseudo-Mask. The best results are in bold.

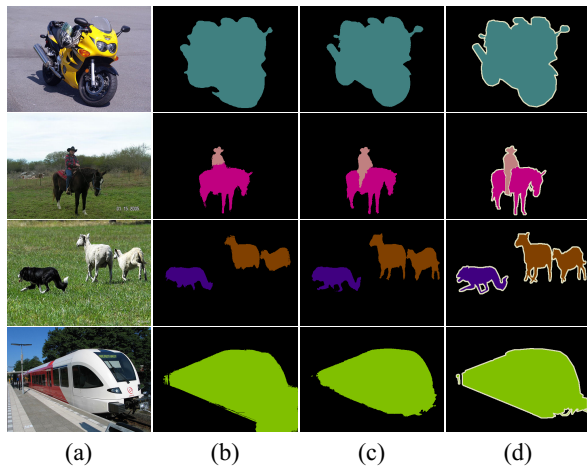


Figure 3: Qualitative results on Pascal VOC 2012 *val* set. (a) Input images. (b) Results from CLIP-ES. (c) Results by our CLIP-ES+VPL. (d) Ground truth. Our method produces more accurate responses and as a plug-and-play method.

with pseudo-labels generated by our method, compared with related works. Our CLIP-ES+VPL achieves state-of-the-art (SOTA) results on VOC, with an mIoU of 78.5% on the *val* set and 77.8% on the *test* set. On the more challenging COCO dataset, our CLIP-ES+VPL surpasses the SOTA results of CLIP-ES and all related works. For CLIP-CPAL, VPL enhances performance (+3.0% mIoU on the COCO validation set). The superior results on both datasets validate the effectiveness of our VPL in accurately capturing semantic features and object structures.

Qualitative Results. We present the qualitative results on the Pascal VOC dataset in Fig. 3. By observation, CLIP-

Methods	VOC		COCO
	Val	Test	Val
Image-level supervision.			
IRN (Ahn, Cho, and Kwak 2019)	63.5	64.8	41.4
OCR-SEAM (Cheng et al. 2023)	67.8	68.4	33.2
ACR (Sun et al. 2023)	71.2	70.9	45.0
ToCo (Ru et al. 2023)	71.1	72.2	42.3
LPCAM (Chen and Sun 2023)	70.1	70.4	45.5
BECO (Rong et al. 2023)	73.7	73.5	45.1
DuPL (Wu et al. 2024a)	73.3	72.8	44.6
Image-level supervision + Language supervision.			
CLIMS (Xie et al. 2022)	70.4	70.0	-
MMCSST (Xu et al. 2023)	72.2	72.2	45.9
WeCLIP (Zhang et al. 2024a)	76.4	77.2	47.1
CLIP-ES (Lin et al. 2023)	73.8	73.9	45.4
+VPL (Ours)	78.5^{↑4.7}	77.8^{↑3.9}	49.2^{↑3.8}
CLIP-CPAL (Tang et al. 2024b)	74.5	74.7	46.8
+VPL (Ours)	79.3^{↑4.8}	79.0^{↑4.3}	49.8^{↑3.0}

Table 2: Evaluation of the mIoU results (%) on PASCAL VOC *val* and *test* sets, and COCO *val* set. The best results for each dataset are in bold.

	Text Pro	Vision Pro	\mathcal{L}_{CE}	RSC(\mathcal{L}_{RSC})	CRF	mIoU(%)
I	✓					63.3
II		✓				68.2
III		✓			✓	71.3
IV		✓	✓			75.5
V		✓	✓	✓		77.2
VI		✓	✓	✓	✓	78.5

Table 3: Ablation study on main components of the proposed framework. The mIoU values are evaluated on the PASCAL VOC 2012 *val* set. In phase 1, Text: Baseline only uses text prototypes. Vision: Vision prototypes for dense localization. CRF: Adoption for post-processing. In phase 2, \mathcal{L}_{CE} : Supervise the network with pseudo-mask. RSC(\mathcal{L}_{RSC}): Align region embedding with the vision prototypes.

ES using text prototypes poorly captures the complete outline of the object as shown in the 2nd rows in Fig. 3. Also, in certain contexts, it has difficulty distinguishing between co-occurring categories, as shown in the 3rd row, such as *train* and *railroad*). Equipped with VPL, the model predicted more accurately on shapes and classes.

Ablation Study

We conduct ablation studies on the PASCAL VOC 2012 *val* set and use CLIP-ES as the Baseline.

Effectiveness of each component. Table 3 presents ablation experiments demonstrating the effectiveness of VPL in two phases. In Experiment I, the model is supervised exclusively using text prototypes, which serve as the primary baseline. In Experiment II, learning vision prototypes to generate GradCAMs enhances performance, with an increase of +4.9%. This indicates that vision prototypes capture fine-grained semantic details better than text prototypes. Subsequently, using CRF to generate pseudo masks boosts performance to 71.3% (Experiment III). In phase 2, we in-

Methods	mIoU(%)	Gap
Baseline _s	61.7	0.81
Baseline _e	63.3	0.76
Modify gap (Liang et al. 2022) +s	62.0	0.77
Modify gap (Liang et al. 2022) +e	63.5	0.71
VPL _s	66.8	0.64
VPL _e	68.2	0.51

Table 4: Comparison of mIoU(%) with different text prompts. “s” denotes a single prompt, while “e” indicates ensemble prompts. “Modify gap” refers to reducing the differences between modalities by adjusting their vector representations (Liang et al. 2022).

τ_I	mIoU(%)	Gap
Baseline	63.3	0.76
0.01	66.9	0.64
0.02	67.3	0.57
0.03	68.2	0.51
0.04	67.8	0.49
0.05	67.2	0.47

Table 5: Impact of temperature τ_I on PASCAL VOC.

roduce cross-entropy loss to supervise the segmentation network (Experiment IV). Specifically, we freeze the CLIP encoder and only train the decoder, resulting in a further increase of +4.2% in performance. In Experiment V, when regional semantic contrast is used as complementary supervision, performance is further improved by +1.7%, indicating its importance in our framework. This module enables the model to focus on the target region, enhancing the matching degree between prototypes and region embeddings.

Impact of temperature. To investigate the impact of the temperature τ_I in learning vision prototypes, Table 5 shows the mIoU and Gap results at different temperatures on the PASCAL VOC 2012 val set. The Gap measures the difference between various prototypes and features, with a lower gap indicating higher similarity. “Baseline” refers to the generation of CAMs using text prototypes. At $\tau_I = 0.03$, the vision prototypes already show an improvement by +4.9% compared to text prototypes. As τ_I increases, the modality gap decreases, enhancing the localization results. This also confirms Proposition 1, which suggests τ_I should be larger than $\tau_T = 0.01$ in CLIP. This indicates that vision and text spaces share common information. By tuning τ_I , we can better utilize this shared information. Thus, we fix $\tau_I = 0.03$ as the optimal value to improve performance.

Prompt Selection. In Table 4, we compare two prompt strategies: using a single prompt like “a photo of a [CLS]” and using an ensemble of multiple prompts such as “a clean origami [CLS].” The placeholder [CLS] represents a category label and its synonyms. Our VPL with the ensemble prompt strategy improves performance over the single prompt by +1.4%. Furthermore, our method outperforms

Φ	0.2	0.3	0.4	0.5	0.6	0.7
mIoU(%)	67.2	67.7	68.2	67.9	67.3	67.9

Table 6: Impact of Threshold Φ on PASCAL VOC.

T_w	1000	1500	2000	2500	3000	3500
mIoU(%)	65.9	66.6	67.5	67.8	68.2	68.0

Table 7: Impact of Iteration T_w on PASCAL VOC.




Figure 4: Feature embedding visualizations of (a) our framework without RSC, and (b) our framework on the PASCAL VOC 2012 val set using t-SNE.

both prompts on text prototypes (“Baseline”), demonstrating the effectiveness of vision prototypes.

Analysis of Hyper-parameters. We conduct a hyperparameter analysis for different values, such as the threshold Φ for generating 0-1 seed masks. In Table 7 (b), the threshold Φ is varied from 0.2 to 0.7 with an interval of 0.1. We find that the optimal choice Φ is 0.4. Moreover, we also investigate the number of iterations T_w for gradient descent to learn the vision prototype and find that the highest accuracy of 68.2% is achieved with 3,000 iterations in Table 7 (c).

Effectiveness of RSC. In Table 3, we present the performance improvement results achieved through the regional semantic contrast module. Moreover, we conduct a visual comparison using t-SNE in Fig. 4. The results indicate that after aligning regional features and prototypes, the model is capable of generating more compact clusters with increased inter-cluster separability.

Conclusion

In this work, we focus on improving the capability of CLIP in weakly supervised semantic segmentation. Our theoretical analysis indicates that the modality gap between text and vision spaces obtained in CLIP is inherent, which can impair the performance. To alleviate this, we propose a Vision Prototype Learning (VPL) framework, which learns more representative prototypes in the vision space and generates high-quality segmentation masks. Extensive experiments have demonstrated that our framework achieves state-of-the-art performance and has the potential to segment new classes.

Acknowledgments

This research was supported by an Australian Government Research Training Program (RTP) Scholarship awarded to Zhongxing Xu.

References

- Ahn, J.; Cho, S.; and Kwak, S. 2019. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*.
- Chen, L.; Lei, C.; Li, R.; Li, S.; Zhang, Z.; and Zhang, L. 2023. FPR: False Positive Rectification for Weakly Supervised Semantic Segmentation. In *ICCV*.
- Chen, Z.; Pan, Y.; Ye, Y.; Lu, M.; and Xia, Y. 2024. Each test image deserves a specific prompt: Continual test-time adaptation for 2d medical image segmentation. In *CVPR*.
- Chen, Z.; and Sun, Q. 2023. Extracting Class Activation Maps from Non-Discriminative Features as well. In *CVPR*.
- Cheng, Z.; Qiao, P.; Li, K.; Li, S.; Wei, P.; Ji, X.; Yuan, L.; Liu, C.; and Chen, J. 2023. Out-of-candidate rectification for weakly supervised semantic segmentation. In *CVPR*.
- Cho, H.; Yoon, S.-H.; Kweon, H.; and Yoon, K.-J. 2024. Finding Meaning in Points: Weakly Supervised Semantic Segmentation for Event Cameras. *ECCV*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *IJCV*.
- Hariharan, B.; Arbeláez, P.; Bourdev, L.; Maji, S.; and Malik, J. 2011. Semantic contours from inverse detectors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 991–998.
- Hu, M.; Xia, P.; Wang, L.; Yan, S.; Tang, F.; Xu, Z.; Luo, Y.; Song, K.; Leitner, J.; Cheng, X.; et al. 2025. Ophnet: A large-scale video benchmark for ophthalmic surgical workflow understanding. In *ECCV*.
- Hu, M.; Yan, S.; Xia, P.; Tang, F.; Li, W.; Duan, P.; Zhang, L.; and Ge, Z. 2024a. Diffusion Model Driven Test-Time Image Adaptation for Robust Skin Lesion Classification. *arXiv preprint arXiv:2405.11289*.
- Hu, M.; Yuan, K.; Shen, Y.; Tang, F.; Xu, X.; Zhou, L.; Li, W.; Chen, Y.; Xu, Z.; Peng, Z.; et al. 2024b. OphCLIP: Hierarchical Retrieval-Augmented Learning for Ophthalmic Surgical Video-Language Pretraining. *arXiv preprint arXiv:2411.15421*.
- Hu, X.; Zhang, K.; Xia, L.; Chen, A.; Luo, J.; Sun, Y.; Wang, K.; Qiao, N.; Zeng, X.; Sun, M.; et al. 2024c. ReCLIP: Refine contrastive language image pre-training with source free domain adaptation. In *WACV*.
- Krähenbühl, P.; and Koltun, V. 2011. Efficient inference in fully connected crfs with gaussian edge potentials. *NeurIPS*.
- Kweon, H.; Yoon, S.-H.; and Yoon, K.-J. 2023. Weakly Supervised Semantic Segmentation via Adversarial Learning of Classifier and Reconstructor. In *CVPR*.
- Lee, J.; Kim, E.; and Yoon, S. 2021. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *CVPR*.
- Lee, J.; Yi, J.; Shin, C.; and Yoon, S. 2021. Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In *CVPR*.
- Li, J.; Jie, Z.; Wang, X.; Wei, X.; and Ma, L. 2022. Expansion and shrinkage of localization for weakly-supervised semantic segmentation. *NeurIPS*.
- Liang, V. W.; Zhang, Y.; Kwon, Y.; Yeung, S.; and Zou, J. Y. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *NeurIPS*.
- Lin, D.; Dai, J.; Jia, J.; He, K.; and Sun, J. 2016. Scribble-sup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*.
- 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*.
- Lin, Y.; Chen, M.; Wang, W.; Wu, B.; Li, K.; Lin, B.; Liu, H.; and He, X. 2023. Clip is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *CVPR*.
- Liu, F.; Liu, Y.; Kong, Y.; Xu, K.; Zhang, L.; Yin, B.; Hancke, G.; and Lau, R. 2023. Referring image segmentation using text supervision. In *ICCV*.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*.
- Ouali, Y.; Bulat, A.; Matinez, B.; and Tzimiropoulos, G. 2023. Black box few-shot adaptation for vision-language models. In *ICCV*.
- Peng, Z.; Wang, G.; Xie, L.; Jiang, D.; Shen, W.; and Tian, Q. 2023. USAGE: A Unified Seed Area Generation Paradigm for Weakly Supervised Semantic Segmentation. In *ICCV*.
- Qian, Q.; Xu, Y.; and Hu, J. 2023. Intra-Modal Proxy Learning for Zero-Shot Visual Categorization with CLIP. *NeurIPS*.
- 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*, 8748–8763. PMLR.
- Rong, S.; Tu, B.; Wang, Z.; and Li, J. 2023. Boundary-Enhanced Co-Training for Weakly Supervised Semantic Segmentation. In *CVPR*.
- Ru, L.; Zheng, H.; Zhan, Y.; and Du, B. 2023. Token contrast for weakly-supervised semantic segmentation. In *CVPR*.
- Sun, W.; Zhang, Y.; Qin, Z.; Liu, Z.; Cheng, L.; Wang, F.; Zhong, Y.; and Barnes, N. 2023. All-pairs Consistency Learning for Weakly Supervised Semantic Segmentation. In *ICCV*.
- Tang, F.; Trinh, M.; Duong, A.; Ly, A.; Stapleton, F.; Chen, Z.; Ge, Z.; and Razzak, I. 2024a. Discriminating retinal microvascular and neuronal differences related to migraines: Deep Learning based Cross-sectional Study. *arXiv preprint arXiv:2408.07293*.
- Tang, F.; Xu, Z.; Huang, Q.; Wang, J.; Hou, X.; Su, J.; and Liu, J. 2023. DuAT: Dual-aggregation transformer network for medical image segmentation. In *PRCV*.

- Tang, F.; Xu, Z.; Qu, Z.; Feng, W.; Jiang, X.; and Ge, Z. 2024b. Hunting Attributes: Context Prototype-Aware Learning for Weakly Supervised Semantic Segmentation. In *CVPR*.
- Trinh, M.; Tang, F.; Ly, A.; Duong, A.; Stapleton, F.; Ge, Z.; and Razzak, I. 2024. Sight for sore heads—using cnns to diagnose migraines. *Investigative Ophthalmology & Visual Science*.
- Udandarao, V.; Gupta, A.; and Albanie, S. 2023. Sus-x: Training-free name-only transfer of vision-language models. In *ICCV*.
- Wan, Z.; Zhang, P.; Wang, Y.; Yong, S.; Stepputtis, S.; Sycara, K.; and Xie, Y. 2024. Sigma: Siamese mamba network for multi-modal semantic segmentation. *arXiv preprint arXiv:2404.04256*.
- Wang, C.; Xu, R.; Xu, S.; Meng, W.; and Zhang, X. 2023. Treating Pseudo-labels Generation as Image Matting for Weakly Supervised Semantic Segmentation. In *ICCV*.
- Wang, J.; Huang, Q.; Tang, F.; Meng, J.; Su, J.; and Song, S. 2022. Stepwise feature fusion: Local guides global. In *MICCAI*. Springer.
- Wang, Y.; Zhang, J.; Kan, M.; Shan, S.; and Chen, X. 2020. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*.
- Wei, J.; and Zhang, X. 2024. Dopra: Decoding over-accumulation penalization and re-allocation in specific weighting layer. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7065–7074.
- Wu, F.; He, J.; Yin, Y.; Hao, Y.; Huang, G.; and Cheng, L. 2024a. Masked Collaborative Contrast for Weakly Supervised Semantic Segmentation. In *WACV*.
- Wu, Y.; Li, X.; Dai, S.; Li, J.; Liu, T.; and Xie, S. 2023. Hierarchical semantic contrast for weakly supervised semantic segmentation. In *IJCAI*.
- Wu, Y.; Ye, X.; Yang, K.; Li, J.; and Li, X. 2024b. DuPL: Dual Student with Trustworthy Progressive Learning for Robust Weakly Supervised Semantic Segmentation. In *CVPR*, 3534–3543.
- Xie, J.; Hou, X.; Ye, K.; and Shen, L. 2022. Clims: Cross language image matching for weakly supervised semantic segmentation. In *CVPR*.
- Xiong, X.; Wu, Z.; Tan, S.; Li, W.; Tang, F.; Chen, Y.; Li, S.; Ma, J.; and Li, G. 2024. Sam2-unet: Segment anything 2 makes strong encoder for natural and medical image segmentation. *arXiv preprint arXiv:2408.08870*.
- Xu, L.; Bennamoun, M.; Boussaid, F.; Laga, H.; Ouyang, W.; and Xu, D. 2024a. Mctformer+: Multi-class token transformer for weakly supervised semantic segmentation. *TPAMI*.
- Xu, L.; Ouyang, W.; Bennamoun, M.; Boussaid, F.; and Xu, D. 2022. Multi-class token transformer for weakly supervised semantic segmentation. In *CVPR*.
- Xu, L.; Ouyang, W.; Bennamoun, M.; Boussaid, F.; and Xu, D. 2023. Learning Multi-Modal Class-Specific Tokens for Weakly Supervised Dense Object Localization. In *CVPR*.
- Xu, Z.; Tang, F.; Chen, Z.; Zhou, Z.; Wu, W.; Yang, Y.; Liang, Y.; Jiang, J.; Cai, X.; and Su, J. 2024b. Polyp-Mamba: Polyp Segmentation with Visual Mamba. In *MICCAI*. Springer.
- Yang, X.; and Gong, X. 2024. Foundation Model Assisted Weakly Supervised Semantic Segmentation. In *WACV*.
- Yang, Y.; Chen, H.; Liu, Z.; Lyu, Y.; Zhang, B.; Wu, S.; Wang, Z.; and Ren, K. 2023. Action recognition with multi-stream motion modeling and mutual information maximization. *arXiv preprint arXiv:2306.07576*.
- Yang, Z.; Fu, K.; Duan, M.; Qu, L.; Wang, S.; and Song, Z. 2024. Separate and Conquer: Decoupling Co-occurrence via Decomposition and Representation for Weakly Supervised Semantic Segmentation. In *CVPR*.
- Ye, Y.; Xie, Y.; Zhang, J.; Chen, Z.; Wu, Q.; and Xia, Y. 2024. Continual self-supervised learning: Towards universal multi-modal medical data representation learning. In *CVPR*.
- Yoon, S.-H.; Kweon, H.; Cho, J.; Kim, S.; and Yoon, K.-J. 2022. Adversarial erasing framework via triplet with gated pyramid pooling layer for weakly supervised semantic segmentation. In *ECCV*.
- Yuan, X.; Shen, C.; Yan, S.; Zhang, X.; Xie, L.; Wang, W.; Guan, R.; Wang, Y.; and Ye, J. 2024. Instance-adaptive Zero-shot Chain-of-Thought Prompting. *arXiv preprint arXiv:2409.20441*.
- Zhang, B.; Yu, S.; Wei, Y.; Zhao, Y.; and Xiao, J. 2024a. Frozen CLIP: A Strong Backbone for Weakly Supervised Semantic Segmentation. In *CVPR*.
- Zhang, P.; Wang, Y.; Liu, Y.; Tu, Z.; and Lu, H. 2024b. Magic tokens: Select diverse tokens for multi-modal object re-identification. In *CVPR*.
- Zhang, X.; Quan, Y.; Gu, C.; Shen, C.; Yuan, X.; Yan, S.; Cheng, H.; Wu, K.; and Ye, J. 2024c. Seeing Clearly by Layer Two: Enhancing Attention Heads to Alleviate Hallucination in LVLMS. *arXiv preprint arXiv:2411.09968*.
- Zhang, X.; Shen, C.; Yuan, X.; Yan, S.; Xie, L.; Wang, W.; Gu, C.; Tang, H.; and Ye, J. 2024d. From Redundancy to Relevance: Enhancing Explainability in Multimodal Large Language Models. *arXiv preprint arXiv:2406.06579*.
- Zhao, X.; Tang, F.; Wang, X.; and Xiao, J. 2024a. SFC: Shared Feature Calibration in Weakly Supervised Semantic Segmentation. In *AAAI*.
- Zhao, X.; Yang, Z.; Dai, T.; Zhang, B.; and Xiao, J. 2024b. PSDPM: Prototype-based Secondary Discriminative Pixels Mining for Weakly Supervised Semantic Segmentation. In *CVPR*, 3437–3446.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *CVPR*.
- Zhou, T.; Zhang, M.; Zhao, F.; and Li, J. 2022. Regional semantic contrast and aggregation for weakly supervised semantic segmentation. In *CVPR*.