

Foundation Model Assisted Weakly Supervised Semantic Segmentation

Xiaobo Yang
 Zhejiang University
 Hangzhou, China
 hal_42@zju.edu.cn

Xiaojin Gong
 Zhejiang University
 Hangzhou, China
 gongxj@zju.edu.cn

Abstract

This work aims to leverage pre-trained foundation models, such as contrastive language-image pre-training (CLIP) and segment anything model (SAM), to address weakly supervised semantic segmentation (WSSS) using image-level labels. To this end, we propose a coarse-to-fine framework based on CLIP and SAM for generating high-quality segmentation seeds. Specifically, we construct an image classification task and a seed segmentation task, which are jointly performed by CLIP with frozen weights and two sets of learnable task-specific prompts. A SAM-based seeding (SAMS) module is designed and applied to each task to produce either coarse or fine seed maps. Moreover, we design a multi-label contrastive loss supervised by image-level labels and a CAM activation loss supervised by the generated coarse seed map. These losses are used to learn the prompts, which are the only parts need to be learned in our framework. Once the prompts are learned, we input each image along with the learned segmentation-specific prompts into CLIP and the SAMS module to produce high-quality segmentation seeds. These seeds serve as pseudo labels to train an off-the-shelf segmentation network like other two-stage WSSS methods. Experiments show that our method achieves the state-of-the-art performance on PASCAL VOC 2012 and competitive results on MS COCO 2014. Our code will be released upon acceptance.

1. Introduction

Weakly supervised semantic segmentation (WSSS) with image-level labels [1, 38, 41, 48] has attracted a great amount of research interest due to its much lower annotation cost compared to fully supervised counterparts [3, 27, 37, 39]. Mainstream methods address this WSSS problem by obtaining reliable seeds and using them as pseudo labels to train an off-the-shelf segmentation network. Therefore, extensive efforts have been devoted to generating high-quality segmentation seeds. These methods often utilize class activation maps (CAMs) [34, 46] derived from image classification networks as initial seeds. However, the obtained CAMs tend to focus on the most discriminative parts rather than complete objects and may falsely activate the background. To improve them, various techniques such as adversarial erasing [18, 19, 38], saliency guidance [21, 44], affinity learning [1, 13, 41], and contrast learning [11, 48] have been exploited. Despite significant advancements achieved by these methods, the performance of WSSS still lags behind that of FSSS counterparts.

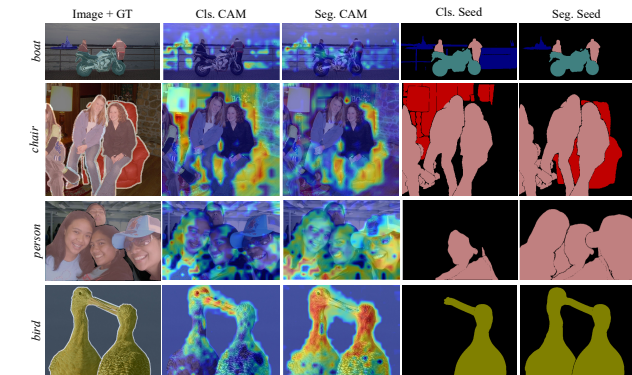


Figure 1. Typical examples of the class activation maps (CAMs) and seed maps produced from the classification (Cls.) and segmentation (Seg.) tasks constructed in our coarse-to-fine framework. The CAMs only demonstrate the maps activated by a target class labelled on the left side, while the seed maps cover all foreground classes presented in the images.

Recently, pre-trained foundation models such as the contrastive language-image pre-training (CLIP) model [28] and the segment anything model (SAM) [16] have also been leveraged to improve the CAMs [15, 24, 40]. While these methods [15, 24, 40] have shown promising results, the way of using foundation models is still in its preliminary stage, leaving ample room for further exploration. For instance, in the adaptation of CLIP to WSSS, both CLIMS [40] and CLIP-ES [24] rely on manually-designed prompts, neglecting the potential benefits of using learnable prompts [47], which have demonstrated effectiveness in other vision tasks. Furthermore, when incorporating SAM into WSSS, a challenge arises as SAM commonly supports point- or box-

wise prompts and produces part/object-level masks without class labels. Previous methods either produce box-wise prompts [36] as an initial step or adopt simple voting schemes [6,15] to select masks. How to effectively integrate SAM into the WSSS framework remains underexplored.

This work attempts to leverage both CLIP and SAM to generate high-quality segmentation seeds for WSSS. Firstly, we introduce prompt learning [47] to adapt CLIP more effectively. A straightforward way to use prompt learning is to replace the manually-designed prompts used in CLIP-ES [24] with contextually learnable prompts trained through an image classification task. However, due to the weak supervision of image-level labels, the *coarse seeds* generated from classification using the learned prompts may still be under-perform, missing object parts or including background regions, as shown in Figure 1. To improve them, we additionally construct a segmentation task and learn a set of segmentation-specific prompts relying on the coarse seeds. This joint framework for classification and segmentation refine the initial segmentation seeds in a progressive manner, moving from a coarse to a fine level. Finally, the *fine seeds* produced by segmentation-specific prompts are used to train the final segmentation network.

In both classification or segmentation tasks of our coarse-to-fine framework, we keep the CLIP model frozen while only learning task-specific prompts. A multi-label contrastive loss relying on image-level labels and a CAM activation loss relying on the generated coarse seeds are designed to train the classification and segmentation tasks jointly. In addition, we employ Softmax-GradCAM [24] to produce the class activation maps in each task stream. Then, instead of taking dense CRF [17] or AffinityNet [1] like previous WSSS methods to refine the CAMs, we design a SAM-based seeding module for refinement, which takes advantage of class-agnostic part/object-level masks produced by SAM to generate seed maps. The CLIP-based coarse-to-fine framework, together with the SAM-based seeding module and the designed losses, empower us to generate high-quality segmentation seeds.

In summary, the contributions of our work are as follows:

- We propose a coarse-to-fine framework that leverages CLIP to jointly perform classification and segmentation tasks. The coarse seed map yielded from classification guides the generation of the fine seed map from segmentation. Moreover, we learn two sets of task-specific prompts, enabling the frozen CLIP model to be adapted to our two tasks more effectively than manually defined prompts.
- We design a SAM-based seeding module that can be applied at the seed generation stage to produce high-quality seeds from CAMs and at the second stage to refine final segmentation results. In contrast to dense CRF [17] or AffinityNet [1] that are extensively used in

previous WSSS methods, our module generates better refinement results while being more efficient.

- Extensive experiments show that our method achieves state-of-the-art on PASCAL VOC 2012 [12] and very competitive results on MS COCO 2014 [23].

2. Related Work

2.1. Weakly Supervised Semantic Segmentation

WSSS methods can be roughly grouped into one-stage and two-stage techniques. The one-stage methods [2,32,45] train a segmentation network end-to-end using image-level labels. In contrast, the two-stage methods [1, 18, 30] first generate segmentation seeds and then use them as pseudo labels to train an off-the-shelf segmentation network, achieving better performance. A key problem in the two-stage methods is to generate high-quality seeds. To this end, various methods such as adversarial erasing [18, 19, 38], saliency guidance [21, 44], affinity learning [1, 13, 41], contrast learning [11, 35, 48], and boundary-aware [22, 30] techniques have been proposed to improve the seeds initially generated by CAMs [34, 46]. Recently, a new trend [6, 24] is to take advantage of pre-trained large foundation models. In line with this research direction, our work has explored incorporating CLIP [24] with SAM [16] to generate high-quality seeds.

2.2. Contrastive Language-Image Pre-training

Contrastive Language-Image Pre-training (CLIP) [28] is a vision-language model trained on 400 million image-text pairs collected from the Internet. It demonstrates a strong ability to generalize to unseen data and has been adapted to various downstream vision tasks. In WSSS, CLIMS [40] proposed CLIP-based losses to supervise another network to generate high-quality CAMs. CLIP-ES [24] designed Softmax-GradCAM and class-aware attention-based affinity (CAA) to generate CAMs directly from CLIP. Both methods [24, 40] use manually-designed templates like “*a photo of [CLS]*” or “*a clean origami [CLS]*” as text prompts when leveraging textual information of CLIP. In contrast, learnable textual contexts have shown their effectiveness in achieving better transferability in other vision tasks [29, 47]. Inspired by them, we propose to learn two sets of prompts specific to classification and segmentation tasks in our coarse-to-fine framework, enabling the frozen CLIP model to be adapted to our tasks more effectively.

2.3. Segment Anything Model

The Segment Anything Model (SAM) [16] is a recently released foundation model for image segmentation. Trained with over 1 billion masks on 11 million images, it has gained a strong generalization ability and can be easily adapted to a range of downstream vision tasks. SAM has

also been applied to boost the performance of WSSS very recently. For instance, Chen *et al.* [6] improved the CAM-based pseudo labels by including SAM-generated segments based on overlap ratios. Jiang *et al.* [15] produced pseudo segmentation labels by prompting SAM with local maximum points on the CAMs. Contrastively, we introduce a SAM-based seeding module that incorporates a filter with a preference order of *whole*, *part*, and *subpart* to generate semantic quasi-superpixels from SAM-generated masks, based on which more complete and precise segmentation seeds are generated. A concurrent work to ours is presented in [36], which employed Grounded-DINO [25] to generate box-wise prompts and feed them into SAM to produce segmentation seeds. In contrast to it, our approach incorporates CLIP with SAM for seed generation.

3. Preliminary

CLIP-ES [24] is a WSSS framework that inspires our work. In this section, we briefly introduce three of its components, including text prompt selection, Softmax-GradCAM, and class-aware attention-based affinity (CAA), which are related to or will be used in our work.

Text prompt selection. CLIP-ES proposes a sharpness-based criterion to guide the choice of text prompts. The selection of prompts is made through trial and error, and “*a clean origami [CLS]*”, which has the lowest sharpness, is manually selected as the input prompt. Here, [CLS] is a foreground class label and its synonyms, or a background label such as *ground*, *grass*, *railroad*, etc., which are additionally defined.

Softmax-GradCAM. CLIP-ES introduces the Softmax function into the original GradCAM [34] to generate CAMs that better suppress non-target classes and backgrounds. Specifically, for each image and a text prompt of any one class label, CLIP yields an image embedding \mathbf{f}_I and a text embedding \mathbf{f}_T , and a logit $Y = \cos(\mathbf{f}_I, \mathbf{f}_T)/\tau$ is computed, where $\cos(\cdot, \cdot)$ denotes cosine similarity and τ is a temperature factor. Then, for the image, a Softmax score s^c for class $c \in \mathcal{P}$ is computed as follows:

$$s^c = \frac{\exp(Y^c)}{\sum_{c' \in \mathcal{P} \cup \mathcal{B}} \exp(Y^{c'})}, \quad (1)$$

where $\mathcal{P} \subseteq \mathcal{F}$ denotes the set of classes presented in the image, and \mathcal{F} is the set of foreground classes in a whole dataset. \mathcal{B} denotes the set of background classes manually defined in CLIP-ES.

Then, given a feature map $F \in \mathbb{R}^{K \times H \times W}$ output before the last multi-head self-attention (MHSA) layer of the CLIP image encoder, the activation weight corresponding to class c at k -th channel is calculated by

$$w_k^c = \frac{1}{HW} \sum_{u=1}^H \sum_{v=1}^W \frac{\partial s^c}{\partial F_{k,u,v}}. \quad (2)$$

The CAM of class c is finally obtained by

$$M_{u,v}^c = \text{ReLU} \left(\sum_{k=1}^K w_k^c F_{k,u,v} \right). \quad (3)$$

Class-aware attention-based affinity. CAA is proposed to refine the above-obtained CAMs by leveraging the attention weight $W^{attn} \in \mathbb{R}^{HW \times HW}$ of the last MHSA layer. That is,

$$M^{c*} = B^c \odot A^t \cdot \text{vec}(M^c), \quad (4)$$

where $B^c \in \mathbb{R}^{1 \times HW}$ is a box mask taken from the CAM of class c , \odot is Hadamard product, t is the iteration number, $\text{vec}(\cdot)$ is the vectorization of a matrix, and A is a symmetric affinity matrix obtained from the attention weight. (Please refer to CLIP-ES [24] for more details.)

4. The Proposed Method

In this section, we present a coarse-to-fine framework based on CLIP [28] and SAM [16] to produce high-quality segmentation seeds. For each image, a coarse seed map is generated from an image classification task that is supervised by a multi-label contrastive loss using the image-level labels. Then, a fine seed map is obtained from a seed segmentation task that is supervised by a CAM activation loss using the generated coarse seeds. We leverage a frozen CLIP model while learning task-specific prompts to perform both classification and segmentation tasks. In each task, Softmax-GradCAM and CAA [24] are applied to produce class activation maps, which are fed into a SAM-based seeding module to generate either coarse or fine seeds. Figure 2 illustrates an overview of the entire framework. Once the task-specific prompts are learned, high-quality segmentation seeds will be produced using the individual segmentation stream and further used as pseudo labels to train an off-the-shelf segmentation network.

4.1. Prompt Learning

In contrast to CLIP-ES [24] that manually selects text prompts, we employ prompt learning [47] to adapt CLIP to WSSS more effectively. Furthermore, instead of learning task-share prompts that overlook fine-grained task correlations, we learn two sets of prompts specific to image classification and seed segmentation tasks, respectively.

Classification-specific prompts. In our WSSS framework, the image-level labels are used to learn a multi-label image classification task. Therefore, we first design a set of prompts specific to this task. Considering that the image-level supervision is relatively weak, we choose a unified context strategy that shares the same context among all classes to avoid overfitting. More specifically, the unified context is defined as follows:

$$t^c = [V]_1[V]_2 \cdots [V]_N[CLS^c], \quad (5)$$

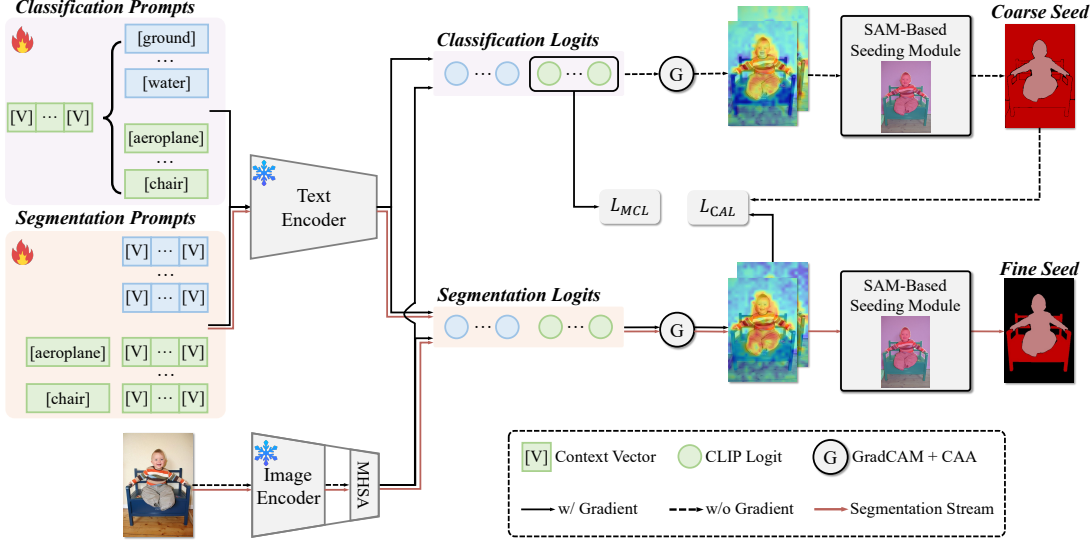


Figure 2. An overview of the proposed coarse-to-fine framework for the generation of high-quality segmentation seeds. It employs a frozen CLIP model with two sets of learnable task-specific prompts to jointly perform image classification and seed segmentation tasks. The logits computed from each set of prompts are used by Softmax-GradCAM to calculate class activation maps (CAMs), which are further refined by class-aware attention-based affinity (CAA). The CAMs obtained from each task are then input into a SAM-based seeding module to generate either a coarse or a fine seed map. Moreover, a multi-label contrastive loss \mathcal{L}_{MCL} and a CAM activation loss \mathcal{L}_{CAL} are designed to supervise the learning of the prompts, which are the only parts that need to be learned in our framework. Once learned, only the segmentation stream, as shown by red lines, is utilized to generate high-quality seeds.

where $[V]_n$ ($n \in \{1, \dots, N\}$) is a vector with the same dimension as the word embeddings of CLIP, and N is the number of context tokens. $c \in \mathcal{F} \cup \mathcal{B}$, in which \mathcal{F} is the set of foreground classes given in a dataset and \mathcal{B} is the set of background classes manually defined in CLIP-ES [24]. This prompt learning technique enables the CLIP model to adapt effectively to our multi-label classification task.

Segmentation-specific prompts. Our framework also includes a seed segmentation task, and therefore, we design a set of segmentation-specific prompts as well. Considering that the supervision of this task is on a pixel level, which is stronger than that of the classification task, we choose a class-specific context strategy in which each class has independent context vectors. That is, for each foreground class $c \in \mathcal{F}$, we define the prompt as

$$t^c = [CLS^c][V]_1^c[V]_2^c \cdots [V]_N^c, \quad (6)$$

and for a background class $b \in \mathcal{B}$ we define the prompt as

$$t^b = [V]_1^b[V]_2^b \cdots [V]_N^b, \quad (7)$$

where $[V]_n^c$ and $[V]_n^b$ have the same dimension as the word embeddings as well.

It is worth noting that, here we prepend the foreground class label to the contexts, which has been validated to be more effective compared to the way used in classification-specific prompts. Additionally, as highlighted by CLIP-ES [24], the regions identified by foreground and background prompts are implicitly mutually exclusive when

employing Softmax-GradCAM. This allows us to eliminate the need for manually defined background class labels in segmentation-specific prompts and still make the background prompts to be fully learned.

4.2. SAM-based Seeding Module

The SAM-based seeding module takes the CAMs as input to produce a seed map. By leveraging the object masks produced by the segment anything model (SAM) [16], we implement this module through three steps: quasi-superpixel generation, quasi-superpixel classification, and seed map generation, which are introduced below.

SAM-based quasi-superpixel generation. When prompted with a regular grid of points, SAM [16] produces ~ 100 high-quality masks per image. However, these masks lack semantic labels, and many of them are overlapped, covering the *whole*, *part*, and *sub-part* of objects. To leverage these masks for generating complete and precise segmentation seeds, we first screen out a set of suitable masks, which we refer to as quasi-superpixels. These quasi-superpixels are selected to preferentially cover whole objects. Additionally, similar to standard non-overlapping superpixels, these quasi-superpixels are chosen to minimize overlap as much as possible.

To this end, we make slight modifications to the mask generation pipeline of SAM as follows: 1) We set a lower confidence threshold $t_{m,whole}$ to retrieve more masks at the *whole* level. 2) During non-maximal suppression (NMS),

we prioritize the selection of *whole* level masks, reducing the suppression of *whole* masks caused by overlaid *part* or *sub-part* masks. 3) We design a simple filter and apply it to the masks retained after NMS to further reduce overlap. The filter follows a preference order of *whole*, *part*, and *sub-part*. That is, it selects *whole* masks first, followed by *part* masks, and then *sub-part* masks. Each subsequent level mask is included only if the occupation ratio between it and the selected masks, which is the ratio between the intersection and the subsequent level mask area, is lower than a threshold t_r . Figure 3 demonstrates typical examples of the masks selected after each of these steps.

SAM-based quasi-superpixel classification. We utilize the refined CAMs to determine the semantic class for each quasi-superpixel. The CAM for a foreground class c is denoted by $M^{c*} \in \mathbb{R}^{H \times W}$, and each quasi-superpixel Q^i is represented by a binary mask of size $H \times W$. The score for Q^i being classified as class c is computed by averaging the activation weights within the quasi-superpixel and then normalizing over all quasi-superpixels. That is,

$$S^c(Q^i) = \text{Norm} \left(\frac{1}{|Q^i|} \sum_{u=1}^H \sum_{v=1}^W Q_{u,v}^i \cdot M_{u,v}^{c*} \right), \quad (8)$$

where $\text{Norm}(\cdot)$ is min-max normalization. The score of being classified as the background is calculated by

$$S^{bg}(Q^i) = (1 - \max_{c \in \mathcal{P}}(S^c(Q^i)))^\alpha, \quad (9)$$

where α is a hyper-parameter empirically set. A larger value of α suppresses the background and results in more foreground superpixels.

Then, the semantic class of quasi-superpixel Q^i is determined by

$$l_i = \arg \max_{j \in \mathcal{P} \cup \{bg\}} (S^j(Q^i)). \quad (10)$$

Note that, unlike prompt learning that considers multiple background classes, here we only consider one background class denoted as bg .

Seed map generation. Finally, we generate a seed map $D \in \mathbb{R}^{H \times W}$ by assigning the class label of each quasi-superpixel to the pixels it contains. If there are some pixels covered by two quasi-superpixels, which is rare after our filtering scheme, the class of the quasi-superpixels at the higher level is assigned. If the levels are the same, the one with the higher classification score is assigned. Additionally, any pixels not covered by any superpixels are treated as background. In contrast to the overlap ratio based voting scheme [6], our *whole*-preference based seed generation is able to retrieve object masks more completely and assign semantic labels more precisely.

4.3. Training Loss

We design a multi-label contrastive loss \mathcal{L}_{MCL} and a CAM activation loss \mathcal{L}_{CAL} to supervise the learning of

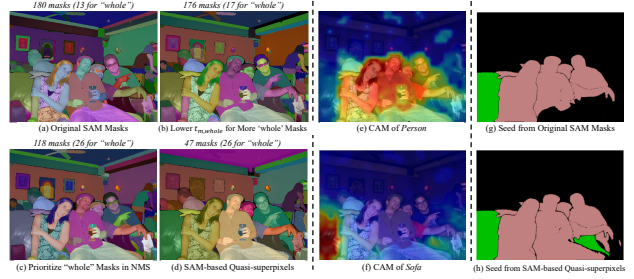


Figure 3. Illustration of the object masks generated by SAM and the resulted seed maps.

prompts. That is, the entire training loss is defined by

$$\mathcal{L} = \mathcal{L}_{MCL} + \mathcal{L}_{CAL}. \quad (11)$$

Multi-label contrastive loss. Although the binary cross entropy loss is commonly used in WSSS to supervise the multi-label classification task, we observe that the values of most CLIP logits for classification fall within the saturation range of the sigmoid function, leading to inefficient training. We hereby employ a multi-label contrastive loss to supervise the classification task, which extends the original contrastive loss of CLIP to the multi-label scenario:

$$\mathcal{L}_{MCL} = \frac{1}{|\mathcal{P}|} \sum_{c \in \mathcal{P}} -\log \frac{\exp(Y_C^c)}{\exp(Y_C^c) + \sum_{c' \in \mathcal{F} - \mathcal{P}} \exp(Y_C^{c'})}, \quad (12)$$

in which Y_C represents a logit for the classification task, \mathcal{P} is the set of classes presented in the image, and $|\cdot|$ is the cardinality of a set. This contrastive loss attracts the image embedding to the text embeddings of foreground classes present in the image while pushing it apart from the text embeddings of absent classes. Note that although the background logits are not considered in this loss, the background prompts still benefit from the unified context shared with the foreground prompts.

CAM activation loss. Although the pixel-wise cross entropy loss is commonly employed for segmentation tasks, it suffers from the low values of Softmax-GradCAM [24] used in our task, leading to ineffective learning. Therefore, the CAM activation loss is designed, which encourages the activation maps yielded from the segmentation task to be closely aligned with the coarse seed map obtained from the classification task. Specifically, let us denote the c -class activation map generated from segmentation as M_S^c and the coarse seed map obtained from classification as D_C . Then, the activation loss for the foreground regions is defined by

$$\mathcal{L}_{CAL}^{fg} = \sum_{c \in \mathcal{P}} \|D_C^c \odot (\max(D_C^c \odot M_S^c) - M_S^c)\|, \quad (13)$$

where $\|\cdot\|$ denotes the L_1 norm. The maximum value is detached to prevent it from being decreased. The loss for the background regions is defined by

$$\mathcal{L}_{CAL}^{bg} = \sum_{c \in \mathcal{P}} \|\overline{D_C^c} \odot M_S^c\|, \quad (14)$$

where $D_C^c = \mathbb{1}(D_C = c)$ is the binary mask for class c and $\overline{D_C^c} = 1 - D_C^c$. Then, the entire CAM activation loss is

$$\mathcal{L}_{CAL} = \frac{1}{|\mathcal{P}| \cdot HW} (\mathcal{L}_{CAL}^{fg} + \mathcal{L}_{CAL}^{bg}). \quad (15)$$

4.4. Final Segmentation

Once the coarse-to-fine framework is learned, we utilize our segmentation stream to generate high-quality seeds. These seeds further serve as pseudo labels to train an off-the-shelf segmentation network, such as DeeplabV3+ [4] or Mask2Former [7], following other WSSS methods. In addition, after obtaining the final segmentation score map, we further apply our SAM-based seeding module but using the segmentation scores for quasi-superpixel classification, to refine the final segmentation results.

5. Experiments

5.1. Datasets and Evaluation Metrics

We evaluate our method on PASCAL VOC 2012 [12] and MS COCO 2014 [23] datasets. PASCAL VOC 2012 [12] provides pixel-wise annotations for 20 object classes and one background class, containing 1464, 1449 and 1456 images for training, validation and testing. Following [14], an augmented set with 10,582 images is used for training. MS COCO 2014 [23] provides annotations for 80 object classes and one background class, containing 82,081 images for training and 40,137 images for validation. We train our network on each dataset using only image-level labels. Following common practice [1, 13, 24, 40], we take the mean intersection-over-union (mIoU) criterion to evaluate performance.

5.2. Implementation Details

The pre-trained ViT-B/16 [10] image encoder and the transformer text encoder of CLIP [28] are adopted and their weights are frozen. The hyper-parameters are set as follows: the number of context tokens N is 16, the threshold t_r is set to 0.3, and α is 0.6. The prompts are learned with a batch size of 16 and 25 epochs on PASCAL VOC and 10 epochs on MS COCO. Other training settings, such as the optimizer, learning rate, etc., are set following CoOp [47].

For the segmentation network trained in the second stage, we utilize either ResNet101-based DeepLabV3+ [4] pre-trained on ImageNet-1k or Swin-L [26]-based Mask2Former [7] pre-trained on ImageNet-21k. When using DeepLabV3+, we employ the multi-scale strategy like other methods while adopting our SAM-based seeding module as the post-processing refinement. When Mask2Former is used, neither the multi-scale strategy nor any post-processing refinement is applied. Additional details and results on more backbone architectures can be found in the supplementary materials.

Table 1. Comparison of the proposed method and its variants. ‘‘Cls.’’ and ‘‘Seg.’’ represent the classification-specific prompts and segmentation-specific prompts learned in our coarse-to-fine framework. ‘‘Post-processing’’ denotes the step used to generate seeds from the CAMs at the seed generation stage.

Models	Prompt Context		Post-Processing		mIoU(%)
	Manual	Cls. Seg.	CRF	SAMS	
CLIP-ES	✓		✓		68.7
\mathcal{M}_0		✓	✓		68.2
\mathcal{M}_1		✓		✓	71.5
\mathcal{M}_2		✓	✓	✓	71.6
\mathcal{M}_3		✓	✓	✓	74.2

Table 2. The performance evaluation of different designs for our segmentation-specific prompts. ‘‘CSC’’denotes the class-specific context strategy. ‘‘Append’’ refers to putting the foreground label after the context tokens, while ‘‘Prepend’’ is putting the label before the contexts. ‘‘w/o bg.’’ denotes that context tokens are not concatenated with background class labels.

Method	mIoU(%)
Unified Context	72.6
CSC + Append + w/ bg.	73.3
CSC + Prepend + w/ bg.	73.6
CSC + Prepend + w/o bg.	74.2

5.3. Ablation Studies

To validate the effectiveness of our designs, we conduct a series of experiments on the PASCAL VOC 2012 *trainaug* set, unless stated otherwise.

Effectiveness of the coarse-to-fine framework. To investigate the effectiveness of our coarse-to-fine framework that includes both classification and segmentation tasks, we compare it to model variants that only involve a classification task. The results are presented in Table 1. By comparing \mathcal{M}_0 to \mathcal{M}_2 or \mathcal{M}_1 to \mathcal{M}_3 , we observe a significant performance boost when incorporating the additional segmentation task.

Effectiveness of the prompt learning. In our framework, we learn two sets of task-specific prompts to adapt CLIP to our two tasks. Table 1 demonstrates that, when only the classification task is considered, the learning of prompts does not improve or even slightly degrades the performance, as seen in the comparison between CLIP-ES [24] and \mathcal{M}_0 . However, the learning of both classification-specific and segmentation-specific prompts achieves a remarkable boost in performance, as shown by \mathcal{M}_2 .

We further investigate the effectiveness of our segmentation-specific prompt design. Specifically, we study the impact of the class-specific context (CSC) strategy versus the unified context strategy, the use of the prepending scheme (*i.e.* put a foreground label before the context tokens) versus the appending scheme (*i.e.* put the foreground label after the contexts), as well as the inclusion

Table 3. The effects of refining the final segmentation results using SAMS or CRF on PASCAL VOC 2012 *val* set.

Backbone	Post-Processing	mIoU(%)
DeepLab V3+	-	75.9
	CRF	75.5
	SAMS	77.3
Mask2Former	-	82.6
	CRF	81.4
	SAMS	82.2

Table 4. The performance evaluation of different ways to generate seeds from the CAMs derived by Softmax-GradCAM. “Baseline” refers to generating seeds directly from the obtained CAMs. “CAA” generate seeds from the CAMs refined by CAA and “Voting” is the method described in [6]. For our SAMS module, “More whole” denotes setting a lower threshold for retrieving more *whole* masks. “NMS Priority” is prioritizing *whole* masks during NMS and “Filter” denotes the additional filtering step.

Method	mIoU(%)	
Baseline	60.7	
+CAA	69.1	
+Voting [6]	69.5	
+CAA+Voting [6]	72.3	
-	71.0	
+SAMS	+ More <i>whole</i>	71.2
	+ NMS Priority	71.5
	+ More <i>whole</i> + NMS Priority	72.0
	+ More <i>whole</i> + NMS Priority + Filter	73.5
+CAA+SAMS	74.2	

or exclusion of manually defined background labels in the background prompts. The comparison results are presented in Table 2, validating that the prompt design we adopt in our framework achieves the best performance.

Effectiveness of the SAM-based seeding module. The SAM-based seeding module (SAMS) can be applied as a post-processing step either at the seed generation stage to generate seeds from the CAMs or at the final segmentation stage to refine the segmentation results. In contrast, many previous methods [24] utilize dense CRF [17] for post-processing. To evaluate the effectiveness of our SAMS, we compare it to CRF at both stages and present the results in Table 1 and Table 3, respectively. Table 1 demonstrates that SAMS significantly outperforms CRF when applied to the seed generation stage. Additionally, Table 3 shows that our SAMS brings a considerable improvement, while CRF degrades the performance when applied to the final segmentation stage using DeepLab V3+.

Within the SAMS module, three modifications are made to the mask generation pipeline of SAM. We here investigate the effectiveness of each modification. Moreover, segmentation seeds can also be generated directly from the CAMs derived from Softmax-GradCAM, or from the CAMs refined by CAA, or from a SAM-based voting scheme simply relied on overlap ratios as described in [6].

Table 5. The performance evaluation of using different classification and segmentation losses. \mathcal{L}_{BCE} stands for the binary cross entropy loss for classification. \mathcal{L}_{CE} represents the pixel-wise cross entropy loss whose probabilities are obtained by inputting CAM activation values into a sigmoid function.

\mathcal{L}_{BCE}	\mathcal{L}_{MCL}	\mathcal{L}_{CE}	\mathcal{L}_{CAL}	mIoU(%)
✓				65.9
	✓			71.5
	✓	✓		52.1
	✓		✓	74.2

Table 6. Quality comparison of the CAMs and segmentation seeds generated by different methods on the PASCAL VOC 2012 *train* set. “Post.” refers to the post-processing techniques used to generate seeds from the obtained CAMs, including dense CRF (CRF), training affinity networks (RW), using class-aware attention-based affinity (CAA) and leveraging SAM.

Methods	Post.	CAM	Seed
AdvCAM _{CVPR21} [19]	RW+CRF	55.6	69.9
MCTformer _{CVPR22} [42]	RW+CRF	61.7	69.1
CLIMS _{CVPR22} [40]	RW+CRF	56.6	70.5
ViT-PCM _{ECCV22} [31]	CRF	67.7	71.4
AdvCAM+W-OoD _{CVPR22} [20]	RW+CRF	59.1	72.1
CLIP-ES _{CVPR23} [24]	CAA+CRF	58.6	75.0
Jiang <i>et al.</i> _{arXiv23} [15]	SAM	47.1	61.9
WeakTr _{arXiv23} [49]	CRF	66.2	68.7
Ours	CAA+SAM	62.6	80.4

Table 4 presents the results obtained using all these different methods and their combinations. We observe that the best performance is achieved by using our proposed SAMS module to generate seeds from the CAMs refined by CAA.

Effectiveness of the training loss. We propose a multi-label contrastive loss and a CAM activation loss to learn the task-specific prompts in our framework. Here, we compare them with the binary cross entropy loss commonly used in the classification task and the pixel-wise cross entropy loss commonly used in the segmentation task. The results are presented in Table 5, demonstrating that both the multi-label contrastive and CAM activation loss outperform their counterparts by a significant margin.

Visualization of segmentation seeds. To make a qualitative comparison, Figure 4 illustrates some segmentation seeds generated by the models defined in Table 1. Compared to CLIP-ES [24] and other model variants, our proposed full model \mathcal{M}_3 generates seeds with more complete regions and precise boundaries for foreground objects. This demonstrates the superior performance of our approach in seed generation for WSSS.

5.4. Comparison to State-of-the-Art

Quality of generated CAMs and segmentation seeds. Table 6 compares the quality of the CAMs and segmentation seeds generated by our model and other methods. The results show that our model produces segmentation seeds

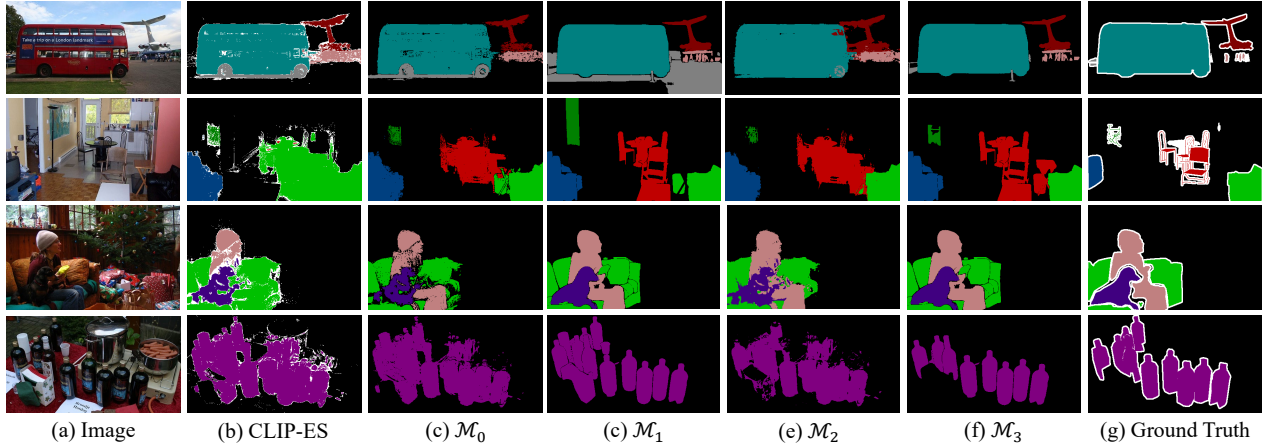


Figure 4. Illustration of segmentation seeds generated by different models.

Table 7. Performance comparison on PASCAL VOC 2012 *val* and *test* sets. The type of supervision used for training is denoted in the "Sup." column, including full supervision (F) and image-level labels (I). The use of CLIP (C), Grounding DINO (D) [25], and SAM (S) are also denoted. The † and ‡ indicate backbone pre-trained on COCO ground-truth [23] or ImageNet-21k [9]. Best WSSS results are marked in bold.

Methods	Backbone	Sup.	Val	Test
DeepLab V3+ _{CVPR18} [4]	R101	F	79.5	79.4
Mask2Former _{CVPR22} [7]	Swin-L‡	F	86.0	86.1
CIAN _{AAAI20} [13]	R101	I	64.3	65.3
AdvCAM _{CVPR21} [19]	R101	I	67.5	67.1
Kweon <i>et al.</i> _{ICCV21} [18]	WR38	I	68.4	68.2
SIPE _{CVPR22} [5]	R101†	I	68.8	69.7
ViT-PCM _{ECCV22} [31]	R101	I	70.3	70.9
CLIMS _{CVPR22} [40]	R50†	I + C	70.4	70.0
AdvCAM+W-OoD _{CVPR22} [20]	WR38	I	70.7	70.1
MCTformer _{CVPR22} [42]	WR38	I	71.9	71.6
ToCo _{CVPR23} [33]	ViT-B‡	I	71.1	72.2
Xu <i>et al.</i> _{CVPR23} [43]	WR38	I + C	72.2	72.2
OCR+MCTformer _{CVPR23} [8]	WR38	I	72.7	72.0
BECO _{CVPR23} [30]	MiT-B2	I	73.7	73.5
CLIP-ES _{CVPR23} [24]	R101†	I + C	73.8	73.9
Jiang <i>et al.</i> _{arXiv23} [15]	R101†	I + S	71.1	72.2
Sun <i>et al.</i> _{arXiv23} [36]	R101†	I + D + S	77.2	77.1
WeakTr _{arXiv23} [49]	DeiT-S	I	74.0	74.1
WeakTr _{arXiv23} [49]	ViT-S‡	I	78.4	79.0
Ours	R101	I + C + S	77.3	76.7
Ours	Swin-L‡	I + C + S	82.6	81.6

better than all previous methods. The high-quality seeds help to train a better segmentation network for WSSS.

Quality of final segmentation results. We further compare the final segmentation results obtained by our method with other state-of-the-art methods on both PASCAL VOC and MS COCO datasets. The results are presented in Table 7 and Table 8. When using the ResNet101 backbone on PASCAL VOC, the proposed method performs comparably to Sun *et al.* [36], which uses a COCO-pretrained backbone. Moreover, both of them outperform other ResNet101-based methods by a significant margin. When the transformer-

Table 8. Performance comparison on MS COCO 2014 *val* set.

Methods	Backbone	Sup.	Val
DeepLab V3+ _{CVPR18} [4]	R101	F	60.4
Mask2Former _{CVPR22} [7]	Swin-L‡	F	66.7
Kweon <i>et al.</i> _{ICCV21} [18]	WR38	I	36.4
AdvCAM _{CVPR21} [19]	R101	I	44.4
SIPE _{CVPR22} [5]	R101	I	40.6
MCTformer _{CVPR22} [42]	WR38	I	42.0
ViT-PCM _{ECCV22} [31]	R101	I	45.0
ToCo _{CVPR23} [33]	ViT-B‡	I	42.3
OCR+MCTformer _{CVPR23} [8]	WR38	I	42.5
BECO _{CVPR23} [30]	R101	I	45.1
CLIP-ES _{CVPR23} [24]	R101	I + C	45.4
Xu <i>et al.</i> _{CVPR23} [43]	WR38	I + C	45.9
WeakTr _{arXiv23} [49]	DeiT-S	I	46.9
WeakTr _{arXiv23} [49]	ViT-S‡	I	50.3
Sun <i>et al.</i> _{arXiv23} [36]	R101	I + D + S	55.6
Ours	R101	I + C + S	48.6
Ours	Swin-L‡	I + C + S	55.4

based backbone is utilized, our methods outperform all other methods. On MS COCO, our method also outperforms all the other methods except Sun *et al.* [36] that leverages Grounding DINO [25] to generate bounding boxes to prompt SAM. In contrast to CLIP [28] that is trained with text-image pairs, Grounding DINO [25] uses fine-grained *phrase-object* pairs for training, leading to richer prior knowledge and better performance.

6. Conclusion

In this work, we have presented a coarse-to-fine framework based on both CLIP and SAM to generate segmentation seeds for WSSS. The proposed task-specific prompt learning, together with the SAM-based seeding module and the designed training losses, enable our approach to generate seeds at high quality. In the proposed framework, the only parts that need to be learned are the two sets of task-specific prompts, and therefore the learning is efficient. Experimental results on PASCAL VOC and MS COCO datasets have validated the effectiveness of our method.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, pages 4981–4990, 2018. 1, 2, 6
- [2] Nikita Araslanov and Stefan Roth. Single-stage semantic segmentation from image labels. In *CVPR*, pages 4253–4262, 2020. 2
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 1–18, 2018. 6, 8
- [5] Qi Chen, Lingxiao Yang, Jianhuang Lai, and Xiaohua Xie. Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation. In *CVPR*, pages 4288–4298, 2022. 8
- [6] Tianle Chen, Zheda Mai, Ruiwen Li, and Wei lun Chao. Segment anything model (sam) enhanced pseudo labels for weakly supervised semantic segmentation. *arXiv preprint arXiv:2305.05803*, 2023. 2, 3, 5, 7
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, pages 1290–1299, 2022. 6, 8
- [8] Zesen Cheng, Pengchong Qiao, Kehan Li, Siheng Li, Pengxu Wei, Xiangyang Ji, Li Yuan, Chang Liu, and Jie Chen. Out-of-candidate rectification for weakly supervised semantic segmentation. In *CVPR*, pages 23673–23684, 2023. 8
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 8
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, pages 1–21, 2021. 6
- [11] Ye Du, Zehua Fu, Qingjie Liu, and Yunhong Wang. Weakly supervised semantic segmentation by pixel-to-prototype contrast. In *CVPR*, pages 4320–4329, 2022. 1, 2
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 2, 6
- [13] Junsong Fan, Zhaoxiang Zhang, Tieniu Tan, Chunfeng Song, and Jun Xiao. Cian: Cross-image affinity net for weakly supervised semantic segmentation. In *AAAI*, pages 10762–10769, 2020. 1, 2, 6, 8
- [14] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, pages 991–998, 2011. 6
- [15] Peng-Tao Jiang and Yuqi Yang. Segment anything is a good pseudo-label generator for weakly supervised semantic segmentation. *arXiv preprint arXiv:2305.01586*, 2023. 1, 2, 3, 7, 8
- [16] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 1, 2, 3, 4
- [17] Philipp Krahenbuhl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NeurIPS*, pages 1–9, 2011. 2, 7
- [18] Hyeokjun Kweon, Sung-Hoon Yoon, Hyeonseong Kim, Daehee Park, and Kuk-Jin Yoon. Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation. In *ICCV*, pages 6994–7003, 2021. 1, 2, 8
- [19] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *CVPR*, pages 4071–4080, 2021. 1, 2, 7, 8
- [20] Jungbeom Lee, Seong Joon Oh, Sangdoon Yun, Junsuk Choe, Eunji Kim, and Sungroh Yoon. Weakly supervised semantic segmentation using out-of-distribution data. In *CVPR*, pages 16897–16906, 2022. 7, 8
- [21] Seungho Lee, Minhyun Lee, Jongwuk Lee, and Hyunjung Shim. Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In *CVPR*, pages 5495–5505, 2021. 1, 2
- [22] Jing Li, Junsong Fan, and Zhaoxiang Zhang. Towards noiseless object contours for weakly supervised semantic segmentation. In *CVPR*, pages 16856–16865, 2022. 2
- [23] 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*, pages 740–755, 2014. 2, 6, 8
- [24] Yuqi Lin, Minghao Chen, Wenxiao Wang, Boxi Wu, Ke Li, Binbin Lin, Haifeng Liu, and Xiaofei He. Clip is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *CVPR*, pages 15305–15314, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [25] 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. 3, 8
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 6
- [27] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 1
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual

- models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 1, 2, 3, 6, 8
- [29] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Densclip: Language-guided dense prediction with context-aware prompting. In *CVPR*, pages 18082–18091, 2022. 2
- [30] Shenghai Rong, Bohai Tu, Zilei Wang, and Junjie Li. Boundary-enhanced co-training for weakly supervised semantic segmentation. In *CVPR*, pages 19574–19584, 2023. 2, 8
- [31] Simone Rossetti, Damiano Zappia, Marta Sanzari, Marco Schaerf, and Fiora Pirri. Max pooling with vision transformers reconciles class and shape in weakly supervised semantic segmentation. In *ECCV*, pages 446–463, 2022. 7, 8
- [32] Lixiang Ru, Yibing Zhan, Baosheng Yu, and Bo Du. Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers. In *CVPR*, pages 16846–16855, 2022. 2
- [33] Lixiang Ru, Heliang Zheng, Yibing Zhan, and Bo Du. Token contrast for weakly-supervised semantic segmentation. In *CVPR*, pages 3093–3102, 2023. 8
- [34] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017. 1, 2, 3
- [35] Shengyang Sun and Xiaojin Gong. Hierarchical semantic contrast for scene-aware video anomaly detection. In *CVPR*, pages 22846–22856, 2023. 2
- [36] Weixuan Sun, Zheyuan Liu, Yanhao Zhang, Yiran Zhong, and Nick Barnes. An alternative to wsss? an empirical study of the segment anything model (sam) on weakly-supervised semantic segmentation problems. *arXiv preprint arXiv:2305.01586*, 2023. 2, 3, 8
- [37] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2021. 1
- [38] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *CVPR*, pages 1568–1576, 2017. 1, 2
- [39] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, pages 1–14, 2021. 1
- [40] Jinheng Xie, Xianxu Hou, Kai Ye, and Linlin Shen. Clims: Cross language image matching for weakly supervised semantic segmentation. In *CVPR*, pages 4483–4492, 2022. 1, 2, 6, 7, 8
- [41] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, Ferdous Sohel, and Dan Xu. Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation. In *ICCV*, pages 6984–6993, 2021. 1, 2
- [42] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Multi-class token transformer for weakly supervised semantic segmentation. In *CVPR*, pages 4310–4319, 2022. 7, 8
- [43] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Learning multi-modal class-specific tokens for weakly supervised dense object localization. In *CVPR*, pages 19596–19605, 2023. 8
- [44] Qi Yao and Xiaojin Gong. Saliency guided self-attention network for weakly and semi-supervised semantic segmentation. *IEEE ACCESS*, 8:14413–14423, 2020. 1, 2
- [45] Bingfeng Zhang, Jimin Xiao, Yunchao Wei, Mingjie Sun, and Kaizhu Huang. Reliability does matter: An end-to-end weakly supervised semantic segmentation approach. In *AAAI*, pages 12765–12772, 2020. 2
- [46] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 1, 2
- [47] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130:2337–2348, 2022. 1, 2, 3, 6
- [48] Tianfei Zhou, Meijie Zhang, Fang Zhao, and Jianwu Li. Regional semantic contrast and aggregation for weakly supervised semantic segmentation. In *CVPR*, pages 4299–4309, 2022. 1, 2
- [49] Lianghui Zhu, Yingyue Li, Jieming Fang, Yan Liu, Hao Xin, Wenyu Liu, and Xinggang Wang. Weaktr: Exploring plain vision transformer for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2304.01184*, 2023. 7, 8