

# Prompting classes: Exploring the Power of Prompt Class Learning in Weakly Supervised Semantic Segmentation


Balamurali Murugesan\*  
ETS Montreal

Rukhshanda Hussain\*  
Jadavpur University

Rajarshi Bhattacharya\*  
Jadavpur University

Ismail Ben Ayed  
ETS Montreal

Jose Dolz  
ETS Montreal



**Figure 1. Impact of the input text prompt on the generation of class activation maps (CAMs).** Employing the ground truth categorical label as [CLS] token (*second column*) does not necessarily result in the best initial CAMs. Furthermore, even though complex techniques to optimize the [CTX] tokens, such as CoOp [76] (*third column*) may improve the CAMs, we have observed that simply modifying the ground truth class in the [CLS] token by a higher correlated synonym leads to improvements in the identified class-related regions (*fourth column*).

## Abstract

Recently, CLIP-based approaches have exhibited remarkable performance on generalization and few-shot learning tasks, fueled by the power of contrastive language-vision pre-training. In particular, prompt tuning has emerged as an effective strategy to adapt the pre-trained language-vision models to downstream tasks by employing task-related textual tokens. Motivated by this progress, in this work we question whether other fundamental problems, such as weakly supervised semantic segmentation (WSSS), can benefit from prompt tuning. Our findings reveal two interesting observations that shed light on the impact of prompt tuning on WSSS. First, modifying only the class token of the text prompt results in a greater impact on the Class Activation Map (CAM), compared to arguably more complex strategies that optimize the context. And second, the class token associated with the image ground truth does

not necessarily correspond to the category that yields the best CAM. Motivated by these observations, we introduce a novel approach based on a **Pr**oMpt **c**lAss **l**Earning (**POLE**) strategy. Through extensive experiments we demonstrate that our simple, yet efficient approach achieves SOTA performance in a well-known WSSS benchmark. These results highlight not only the benefits of language-vision models in WSSS but also the potential of prompt learning for this problem. The code is available at [code link](#)

## 1. Introduction

Image semantic segmentation is a fundamental problem in computer vision, as it serves as a precursor of many tasks, such as medical image analysis or autonomous driving. Fu-

\* Authors contributed equally to this article. R. Hussain and R. Bhattacharya did this work as part of a research internship at ETS Montreal.

eled by the advances in deep learning, semantic segmentation has experienced a tremendous progress. Nevertheless, obtaining precise pixel-wise annotations is a labor-intensive and time-consuming task.

To alleviate the annotation burden, weakly supervised semantic segmentation (WSSS) has emerged as an appealing alternative, where labels typically come in the form of image tags [15, 21, 28, 31], bounding boxes [26, 55], scribbles [41, 57] or global constraints [25, 46], among others. In particular, image-level WSSS has received significant attention, as it offers a cost-effective alternative to pixel-level annotations (e.g., 20 seconds reported in [5]). Under this setting, WSSS commonly leverages class activation maps (CAMs) [75] obtained from image classification networks to localize objects. Specifically, these maps are later used as pixel-wise pseudo-labels to train a segmentation model, mimicking full supervision. However, CAMs tend to highlight discriminative regions, while ignoring other useful cues, which results in suboptimal pseudo-labels that do not cover the whole extent of the target objects. Narrowing down the existing gap between classification and segmentation tasks is therefore crucial for the progress of WSSS models. To solve this issue, existing approaches intend to complete generated CAMs by forcing the network to focus on more non-discriminative regions, which can be achieved by region mining strategies [21, 30], or integrating iterative processes [3]. Despite employing complex CAM refinement strategies, sometimes involving multiple training steps, existing approaches still exhibit suboptimal performance in terms of both completeness of the target objects and segmentation accuracy.

This motivates the exploration of complementary learning strategies that can further improve the segmentation performance of these models. Vision-language pre-training (VLP) models, such as the recently introduced Contrastive Language-Image Pre-training (CLIP) [47] strategy, have the potential to bring WSSS approaches to the next level, as it can associate much wider visual concepts in an image with their corresponding text labels in an open-world scenario. This contrasts with standard WSSS settings, where the fixed set of predetermined object categories limits the quality of generated CAMs due to unnecessary background activations from class-related background pixels. For example, CLIMS [66] exposed these issues showing that background pixels related to the class ‘railroad’ contributed to the prediction of the CAM associated to the category ‘train’, leading to over-segmented CAMs.

With the rise of these powerful vision-language pre-training models, recent evidence has highlighted the importance of their text input, typically referred to as *prompt*, in adapting these models to downstream tasks and datasets. For instance, Zhou *et al.* [77] empirically demonstrated that the use of ‘a [CLS]’ or ‘a photo of a [CLS]’ as a prompt led to substantial differences in the classification performance

of the model. Following these findings, recent literature has focused on tuning the context of these prompts, typically as continuous learnable vectors [24, 76, 77]. Despite its potential importance, the impact of modifying the [CLS] token has been largely overlooked in the context of prompt learning. Additionally, while prompt learning has shown promising results in fine-tuning and classification tasks such as zero-shot image recognition, its effectiveness on other visual recognition problems is not well-understood.

Based on these observations, we explore in this work how vision-language pre-training can be further leveraged to improve the performance of WSSS models. In particular, we want to address the following questions: ① *Is prompt learning useful in weakly supervised segmentation?*, ② *Which parts of the prompt have a greater impact on the generated CAMs?* ③ *Can we devise a simple yet effective alternative to improve the segmentation performance under the weakly-supervised learning paradigm?*

**Our contributions** can be summarized as follows:

- We provide empirical evidence that modifying the input prompt in VLP models has a direct impact on the generated CAMs in a weakly supervised segmentation scenario, which in turn affects the performance of the segmentation network.
- More interestingly, our findings reveal that replacing the [CLS] token in the input prompt has a greater impact on the performance than modifying the prompt context, which contrasts with recent observations in classification problems (See Fig. 1). Furthermore, the [CLS] token associated with the actual image ground truth does not necessarily correspond to the category that yields the best CAM, and the performance varies considerably across closely related categories. These insights shed light on the importance of careful prompt design in optimizing the performance of segmentation models trained under the weakly-supervised paradigm.
- Based on these observations, we propose a simple yet efficient strategy to leverage language driven models in the challenging task of WSSS. The resulting model, based on a **Pr**Ompt **c**Lass **I**Earning (**POLE**) approach, learns the category name that produces the highest correlation between the image and a corresponding text prompt, and uses it to further leverage the segmentation performance.
- Following the literature, we conduct extensive experiments on PASCAL VOC 2012 to well demonstrate the superiority of our method over other state-of-the-art methods for WSSS.

## 2. Related Work

**Weakly supervised semantic segmentation.** Due to its low annotation cost, WSSS based on image-level labels has

gained increasing popularity. These methods rely on class activation maps (CAMs) to identify target object regions by discovering informative pixels for the classification task. As discovered regions are typically highly discriminative and fail to cover the whole context of the target objects, recent literature focuses on generating high-quality CAMs by refining initial estimations from simple models. A common strategy is to mine or erase regions at either image [30, 61, 73] or features level [21, 31], and can be seen as a way of preventing a classifier from focusing exclusively in highly discriminative areas. Other works have instead exploited sub-categories dependencies [6], cross-image semantics [16, 56], attention mechanisms [63], equivariant constraints [45, 60] and pairwise semantic affinities [3, 59]. Furthermore, additional supervision, such as saliency maps can be also integrated to provide additional hints about the location of the target object [23, 35]. More recently, visual transformers (ViT) [12] have been also leveraged to improve original CAMs [39, 69], demonstrating superior performance than their CNNs counterparts.

**Contrastive Language-Image Pre-training (CLIP) based semantic segmentation.** Very recently, large-scale VLP models, such as CLIP [47], have demonstrated to improve significantly the performance of vision models on classic recognition tasks, such as zero-shot object detection [18], few-shot learning [22] and zero-shot semantic segmentation [37, 78]. Closely related to our work, CLIMS [66] integrates CLIP in the context of weakly-supervised segmentation, which enhances the initial CAMs by highlighting more comprehensive object regions, while suppressing closely-related background areas. Inspired by the improvement observed in the robustness and generability of visual recognition models driven by language assistance, our work delves deeper into understudied factors, particularly in the weakly-supervised scenario. We stress that our work is different from [66]. In particular, CLIMS [66] proposes to leverage standard CLIP in WSSS, whereas we further explore the effect of the given prompt on this task. As we will show in our empirical validation, properly designing the input prompt results in significant improvements over the standard text prompts. More surprisingly, *using the class ground truth as categorical name in the input text prompt does not necessarily yields the best segmentation results.*

**Prompt learning** in visual recognition problems is a rapidly growing research direction, whose popularity stems from the promising results observed in NLP tasks [36, 40, 42, 48]. For example, recent works in prompt learning [13, 44, 49, 58, 76, 77] have achieved promising results on several vision-language tasks, notably in classification. In addition to tackle a different task, i.e., classification *vs.* weakly supervised segmentation, the main differences with our work is that these approaches mostly study prompt learning from a context perspective. In particular, existing literature considers the class token [CLS] as a fixed word embedding, while

optimizing the context [13, 49, 58, 76, 77] or attributes [44]. In most cases, the context tokens are represented by learnable continuous vectors, which yield to text embeddings lacking semantic knowledge [77]. In contrast, our approach performs a selection on a finite set of potential synonyms, which facilitates both the search and the interpretation of the selected token.

## 3. Methodology

### 3.1. Problem setting.

Let us denote  $\mathcal{D} = \{(\mathbf{X}_i, \mathbf{y}_i)\}_{i=1}^N$  as a weakly labeled dataset, where  $\mathbf{X}_i \in \mathbb{R}^{\Omega_i}$  is an input image,  $\Omega_i$  denotes its spatial dimensionality,  $\mathbf{y}_i \in \{0, 1\}^K$  its associated one-hot encoded image label<sup>1</sup>, and  $K$  indicates the number of categories. Thus, the goal of weakly supervised semantic segmentation is to provide pixel-wise predictions from an input image  $\mathbf{X}_i$  given its corresponding image-level label  $\mathbf{y}_i$ .

### 3.2. Our framework.

**Preliminaries: Class Activation Maps.** We first revisit the generation of class activation maps (CAM) from the image-level labels, a popular strategy in WSSS. Let us first define a feature extractor  $f_\theta(\cdot)$ , which can be represented by a deep neural network parameterized by  $\theta$ . Thus, for a given image  $\mathbf{X}$ , the feature extractor provides a representation  $\mathbf{Z} \in \mathbb{R}^{C \times H' \times W'}$ , where  $C$  is the number of channels and  $H'$  and  $W'$  represent the dimensionality of the feature map. To provide CAMs, a global average pooling (GAP) layer, followed by a  $1 \times 1$  convolutional layer  $\mathbf{W} \in \mathbb{R}^{C \times K}$  is applied to the learned features  $\mathbf{Z}$  from an image  $\mathbf{X}$ . Then, the resulting logits are mapped into probabilities  $\hat{\mathbf{y}} \in [0, 1]^K$  by applying a sigmoid function. To train the neural network, we follow the literature [66, 67, 69] and use the multi-label soft-margin loss as the classification function:

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) = -\frac{1}{K} \sum_{k=1}^K (y^k \log \hat{y}^k + (1 - y_k) \cdot \log(1 - \hat{y}_k)). \quad (1)$$

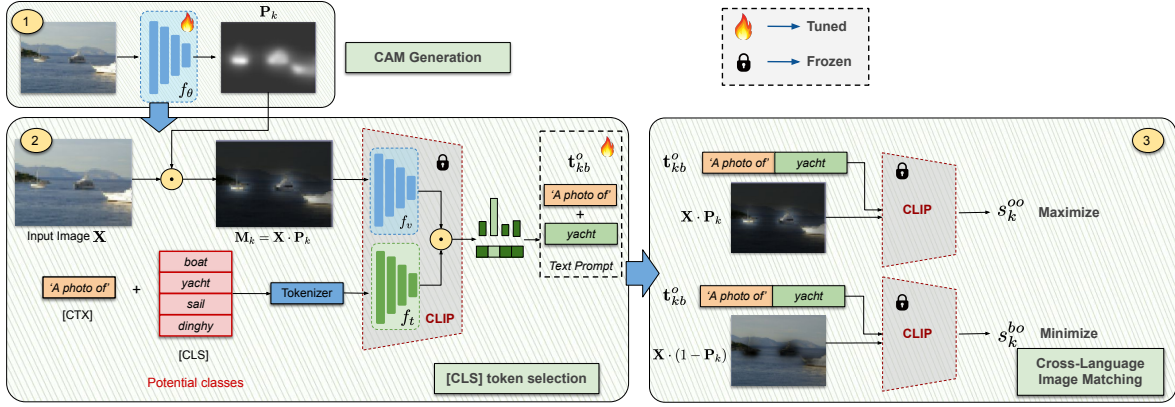
Once the backbone network is trained, the initial CAMs can be obtained as follows:

$$\mathbf{P}_k(h, w) = \mathbf{W}_k^\top \mathbf{Z}(h, w), \quad (2)$$

where  $\mathbf{P}_k$  is the activation map for a given category  $k$ .

**Learning objectives.** The framework used for this work shares the same overall structure as the recent CLIMS [66], as it represents the first weakly-supervised segmentation approach integrating CLIP text embeddings. Nevertheless, we stress that in our work we study how we can leverage prompt learning to further improve the performance

<sup>1</sup>In PascalVOC, multiple classes can be present in the same image, where  $\mathbf{y}$  becomes a multi-class one-hot encoded vector.



**Figure 2. Proposed Weakly Supervised Segmentation approach.** 1) Class activation maps are generated for an input image  $\mathbf{X}$ . 2) CLIP pre-trained visual and text encoders ( $f_\theta$  and  $f_\theta$ ) are leveraged to find the category name [CLS] presenting the highest correlation with the image  $\mathbf{M}_k$ , the result of multiplying the input image  $\mathbf{X}$  and its corresponding CAM  $\mathbf{P}_k$ . 3) With the [CLS] token selected, we generate the input text prompt  $\mathbf{t}_{k,b}^o$  to the Cross-Language Image Matching (CLIMS) learning framework.

of weakly supervised segmentation models. In particular, the standard GAP layer employed to generate CAMs is replaced by a sigmoid function  $\sigma(\cdot)$ , resulting in  $\mathbf{P}_k(h, w) = \sigma(\mathbf{W}_k^T \mathbf{Z}(h, w))$ . The pre-trained CLIP model [47] uses a visual and a text encoder which we denote as:  $f_v$  and  $f_t$ , respectively. Instead of passing the raw input image  $\mathbf{X}$  to the CLIP image encoder, it is multiplied by  $\mathbf{P}_k$ , with the goal of focusing only on the class highlighted by its corresponding CAM. Moreover, to avoid interferences from related background regions, the input image is also multiplied by  $(1 - \mathbf{P}_k)$ . Hence, we can create two different visual embeddings: the embedding of the target category ( $\mathbf{v}_k^{io}$ ), and its background ( $\mathbf{v}_k^{ib}$ ), which are formally given by:

$$\mathbf{v}_k^{io} = f_v(\mathbf{X} \cdot \mathbf{P}_k), \quad \mathbf{v}_k^{ib} = f_v(\mathbf{X} \cdot (1 - \mathbf{P}_k)). \quad (3)$$

Now, for all the potential object classes  $k$  and their corresponding text inputs  $\mathbf{t}_k^o$ , we obtain their text embeddings, which are referred to as  $\mathbf{v}_k^{to}$ :

$$\mathbf{v}_k^{to} = f_t(\mathbf{t}_k^o). \quad (4)$$

Note that to generate these embeddings, we just need to provide the different text inputs to the trained CLIP text encoder  $f_t(\cdot)$ . Following the reasoning behind the training of CLIP, the foreground image embedding  $\mathbf{v}_k^{io}$  should be highly correlated to the text embedding  $\mathbf{v}_k^{to}$  of that particular class. In contrast, the background image embedding  $\mathbf{v}_k^{ib}$  should have a much lower correlation with the object classes. This can be modeled by using the following objective function:

$$\mathcal{L}_{Cont} = -\alpha \sum_{k=1}^K y_k \cdot \log(s_k^{oo}) - \beta \sum_{k=1}^K y_k \cdot \log(1 - s_k^{bo}), \quad (5)$$

where  $s_k^{oo} = \text{sim}(\mathbf{v}_k^{io}, \mathbf{v}_k^{to})$  and  $s_k^{bo} = \text{sim}(\mathbf{v}_k^{ib}, \mathbf{v}_k^{to})$  represent the *object-to-object* and *background-to-object* simi-

larities between visual and text embeddings, computed as a cosine similarity. Both terms in Eq. 5 act together to ensure that the activation map  $\mathbf{P}_k$  covers the maximum possible region of the target, while excluding related background.

### 3.3. Category and image-driven prompt generation

**Finding potential category related embeddings.** To generate the object text representation  $\mathbf{v}_k^{to}$  for a class  $k$ , the standard input prompt given to the text encoder has the following format: a context token [CTX] followed by the class name token [CLS] and ended by a punctuator ('.'). While the literature in prompt learning for large-scale visual language pre-trained models focuses on learning the context [CTX] [76], CLIMS [66] uses a fixed prompt, where the [CLS] token corresponds to the categorical label of the image. Contrary to these works, we hypothesize that modifying the input text prompts by optimizing only the [CLS] token has a greater impact on the generated CAMs. Indeed, as we will show empirically in the results section, using the ground truth class as a [CLS] token does not necessarily always results in the best segmentation performance.

Let us suppose that we take an input image  $\mathbf{X}$  with its corresponding image class label  $\mathbf{y}$ , which indicates the  $k$  categories present on the image. For each category  $k$  in  $\mathbf{y}$ , we obtain a set of similar words, in terms of closeness in the semantic space, using chatGPT [1]. More concretely, we provide the following query as input to chatGPT "Give me  $m$  semantically similar words for [CLS] and also print the cosine similarity scores based on CLIP model", where [CLS] is a class name. This returns a list of  $m$  words along with their similarity scores for that particular [CLS]. This means that for each class [CLS], we can derive a set of  $m$  closest words, denoted as  $\mathcal{S} = [\text{CLS}_1, \text{CLS}_2, \dots, \text{CLS}_m]$ . With this set of related categories, we can create a set of  $m + 1$  potential text prompts  $\mathcal{T} = [\mathbf{t}_{k,0}^o, \mathbf{t}_{k,1}^o, \mathbf{t}_{k,2}^o, \dots, \mathbf{t}_{k,m}^o]$ ,

where  $\mathbf{t}_{k0}^o$  is the text prompt containing the categorical ground truth label for class  $k$ , i.e., [CLS] followed by a fixed [CTX] token, and  $\mathbf{t}_{k1}^o, \dots, \mathbf{t}_{km}^o$  are composed of the fixed [CTX] followed by a variable [CLS] token chosen from set  $\mathcal{S}$ . Now, we can extract an embedding for each of the prompts in  $\mathcal{T}$  from the CLIP text encoder, resulting in  $\mathcal{V} = [\mathbf{v}_{k0}^{to}, \mathbf{v}_{k1}^{to}, \mathbf{v}_{k2}^{to}, \dots, \mathbf{v}_{km}^{to}]$ .

**How to select the best [CLS]?** Given the input image  $\mathbf{X}$  and its generated class activation map  $\mathbf{P}_k$  for class  $k$ , we can obtain an image focusing on discriminative regions for that category by simply doing  $\mathbf{M}_k = \mathbf{X} \cdot \mathbf{P}_k$ . The resulting image,  $\mathbf{M}_k$ , is given to the CLIP image encoder to get a compressed representation of  $\mathbf{X}$ , i.e.,  $\mathbf{v}_k^{io}$ . similar to the steps described in Section 3.2. In contrast, we now obtain a correlation between each of the closest words in set  $\mathcal{S}$  for the class  $k$  and the CAM activated image  $\mathbf{M}_k$ . In particular, this correlation is found by computing a similarity score between the visual and text embeddings:  $\mathbf{v}_k^{io}$  and every  $\mathbf{v}_{kj}^{to}$  for  $j \in \{0, 1, 2, 3, \dots, m\}$ , given as:

$$\text{sim}(\mathbf{v}_k^{io}, \mathbf{v}_{kj}^{to}) = \frac{\mathbf{v}_k^{io} \cdot \mathbf{v}_{kj}^{to}}{|\mathbf{v}_k^{io}| |\mathbf{v}_{kj}^{to}|}, \quad (6)$$

which generates a vector containing the similarities between the visual encoding and each of the text encodings  $[s_{k0}, s_{k1}, \dots, s_{km}]$ . From this similarity vector, we select the most correlated [CLS] token, which corresponds to the text embedding  $\mathbf{v}_{kj}^{to}$  with the highest similarity with  $\mathbf{v}_k^{io}$ , computed with the argmax operator.

### 3.4. Weakly supervised adaptors

Following the success of adaptors in pre-trained language-vision models for classification tasks [17, 49, 74], we propose to further improve our segmentation network by integrating image and text adaptors. In particular, and similar to [17] in classification, we introduce two MLP layers  $A_v(\cdot)$  and  $A_t(\cdot)$  to transform the embeddings in the image side and text space, respectively, which is formulated as:

$$\mathbf{v}_k^{io*} = \mathbf{r}_v \cdot A_v(\mathbf{v}_k^{io}) + (1 - \mathbf{r}_v) \cdot \mathbf{v}_k^{io} \quad (7)$$

$$\mathbf{v}_k^{to*} = \mathbf{r}_t \cdot A_t(\mathbf{v}_k^{to}) + (1 - \mathbf{r}_t) \cdot \mathbf{v}_k^{to}, \quad (8)$$

with  $A_v(\mathbf{v}_k^{io}) = \text{ReLU}(\mathbf{v}_k^{io} \mathbf{W}_1^v) \mathbf{W}_2^v$  and  $A_t(\mathbf{v}_k^{to}) = \text{ReLU}(\mathbf{v}_k^{to} \mathbf{W}_1^t) \mathbf{W}_2^t$ , where  $\mathbf{W}$  represents the learnable parameters of the MLP layers. Furthermore, the parameters  $\mathbf{r}_v$  and  $\mathbf{r}_t$  are learnable vectors of the same shape as the original embeddings, and are used to selectively suppress or amplify the refinement of each feature through the MLP layers. Note that this contrasts with standard adaptors, i.e., [17, 49, 74], whose balancing weight is a fixed hyperparameter. Our hypothesis is that using a fixed scalar to control the importance of each embedding is suboptimal, as the features refinement process may differ across images as well as depend on the class variation of the dataset.

## 4. Experiments

### 4.1. Experimental Settings

**Dataset and evaluation protocol.** Following CLIMS [66], as well as other recent WSSS works [32], we conduct experiments on the popular PASCAL VOC 2012 [14] benchmark. This dataset contains images with 20 foreground classes, which are split into 1,464 for training, 1,449 for validation and 1,456 for testing. The training set is augmented with 10,582 images and their associated image-level annotations from SBD [19]. To evaluate the performance of the proposed method, we resort to the mean intersection over union (mIoU). Last, while the results reported in the ablation studies are obtained on the training set, the results for the validation and testing sets of PASCAL VOC are obtained from the official evaluation server.

**Implementation Details.** We followed the default settings of CLIMS [66] for training. In particular, input images are randomly rescaled and then augmented by random cropping to  $512 \times 512$ , as well as by horizontal flipping. We use SGD as the default optimizer, with a cosine annealing policy for scheduling the learning rate, and a batch size of 16 images. The model is trained for 10 epochs, with an initial learning rate of 0.00025 and a weight decay of 0.0001. We follow [4] to adopt ResNet-50 [20] as backbone network for the generation of initial CAMs. All models are implemented in PyTorch and trained on NVIDIA A100 GPU with 40 GB memory. Furthermore, as the initial CAMs coarsely covers the target object, we further perform a refinement step with IRNet [2], to improve their quality before using them as pseudo ground-truth masks, a common practice in WSSS.

### 4.2. Results

**How effective is prompt learning for weakly supervised segmentation?** In this section we assess whether modifying the input text prompt leads to performance differences, which corresponds to our first question. To do this, we first manually selected two popular prompts: “A photo of [CLS].”, and “An image of [CLS].”, which are employed over all the images of the whole dataset. Note that the original CLIMS [66] used the first prompt, and it is therefore considered as the baseline model. Furthermore, inspired by the recent advances on prompt learning, we also evaluate the impact different strategies, which model the different tokens with learnable continuous vectors: CoOp [77] ([CTX] tokens), *target optimization* baseline in [58] ([CLS] tokens) and a modified version of DeFo [58]. These results, which are reported in Table 1, reveal that the text input, i.e., prompt, of the pre-trained vision-language model plays an important role on the segmentation performance. Indeed, we can observe that depending on the prompt employed, performance differences may diverge up to 3%, particularly on the final generated CAM (*last column*).

**Context vs. category in prompt learning.** Once we have

**Table 1. Does prompt learning improve the performance of weakly supervised segmentation?** Comparison of the quality of initial CAMs and refined pseudo ground-truth masks obtained by different prompt learning strategies (with R50 as backbone), where either [CTX] or [CLS] tokens are modified. Evaluation is reported on the *train* set of PASCAL VOC2012, and refinement of the pseudo-masks is performed using **RW** (IRN [2]). [CTX] and [CLS] are used to indicate which part of the prompt is optimized in each approach. Furthermore, in the approaches optimizing the [CLS] token, 'V' indicates a continuous learnable vector, whereas CLS\* represents the class selected among a set of potential classes.

Method	[CTX]	[CLS]	Prompt	CAMs	+RW
Manual [66]	✓	✗	"A photo of [CLS]."	56.6	70.5
Manual	✓	✗	"An image of [CLS]."	56.5	71.0
CoOp [76]	✓	✗	[V] <sub>1</sub> [V] <sub>2</sub> ...[V] <sub>N</sub> [CLS].	57.6	73.1
DeFo [58]	✓	✓	[V] <sub>1</sub> [V] <sub>2</sub> ...[V] <sub>N</sub> [V <sub>CLS</sub> ]	56.6	73.2
Target optimization	✗	✓	"A photo of [V <sub>CLS</sub> ]."	56.8	73.1
Ours	✗	✓	"A photo of [CLS*]."	<b>58.7</b>	<b>73.6</b>

**Table 2.** Comparison of the quality of initial CAMs and refined pseudo ground-truth masks using **RW** (PSA [4]) on PASCAL VOC2012. The mIoU values here are reported on the *train* set. Backbone denotes the backbone network to generate CAMs. Best approaches (CAM and refined CAM) highlighted in bold.

Method	Backbone	CAMs	+RW
PSA CVPR'2018 [3]	WR38	48.0	61.0
SC-CAM CVPR'2020 [6]	WR38	50.9	63.4
SEAM CVPR'2020 [60]	WR38	55.4	63.6
AdvCAM CVPR'2021 [32]	R50	55.6	68.0
MCTformer CVPR'2022 [69]	ViT	<b>61.7</b>	69.1
SIPE CVPR'2022 [9]	R50	58.6	64.7
RECAM CVPR'2022 [10]	R50	54.8	70.5
AdvCAM+W-OoD CVPR'2022 [33]	R50	59.1	72.1
AFA CVPR'2022 [53]	MiT		68.7
CLIMS CVPR'2022 [66]	R50	56.6	70.5
VWL IJCV'2022 [52]	R101	57.3	71.4
AEFT ECCV'2022 [71]	WR38	56.0	71.0
ViT-PCM ECCV'2022 [50]	ViT-B/16	–	71.4
ESOL NeurIPS'2022 [38]	R50	53.6	68.7
<b>POLE (Ours)</b>	R50	59.0	<b>74.2</b>

observed empirically that modifying input prompts results in performance differences, one question that naturally arises is which component of the prompt must be changed. Table 1 shows that replacing a standard [CTX] token (i.e., 'A photo of') by a similar sequence (i.e., 'An image of') brings 0.5% difference. Nevertheless, if the [CTX] token is optimized for the whole dataset, e.g., CoOp [76], these differences are further increased, with similar results if [CLS] token is optimized as a continuous vector (e.g., DeFo [58] and *Target Optimization*). Last, we can observe that only optimizing the [CLS] prompt, based on a set of pre-defined closely-related categories, actually provides the best performance across all the methods. These findings align with recent observations in Natural Language Inference [43], which suggest that hand-crafted prompts conveying meaningful instructions outperform automatically optimized prompts.

**Comparison to state-of-the-art.** Previous experiments empirically demonstrated that modifying the input prompt can significantly improve the performance of weakly supervised segmentation models. These observations motivated the proposed approach, which we benchmark against state-of-the-art models to show its superiority. Note that in what follows, the proposed approach, POLE, is composed of the [CLS] token selection process and the adaptor with learnable weights (eq. 7). First, Table 2 reports the results of state-of-the-art methods in the generation of pseudo-masks. We can observe that, even compared to very recent models, the proposed approach brings substantial improvements, ranging from 2 to 6%. We interpret that the performance gain observed comes from the highest correlation between the selected category name and the content of the CAM-activated region in the image., which may capture larger discriminant areas of semantic objects. More interestingly, the proposed approach also outperforms recent methods that use additional information, for example in the form of extra saliency maps, which are typically trained on a supervised foreground-background detection dataset. Thus, based on these results we can argue that our approach represents an effective alternative to generate initial CAMs.

While the quality of the initial CAMs was evaluated in the previous section, we now assess their impact on the semantic segmentation task. In particular, Table 3 reports the segmentation performance of the proposed approach compared to state-of-the-art methods in the validation and testing sets of PASCAL VOC2012 dataset. Compared to approaches that resort to the same supervision, our method provides very satisfactory results, ranking second if all the approaches are considered. Nevertheless, it has been found recently that vision transformers provide much better quality CAMs than conventional convolutional neural networks. Thus, if we just consider the approaches that leverage CNNs for the CAM generation, our approach achieves the best performance, with improvement gains ranging from 0.7% to 3% compared to very recent methods (e.g., RECAM, SIPE or CLIMS [66]). Furthermore, even benchmarking POLE against recent approaches that use additional supervision, e.g., saliency maps, it yields very competitive performance.


**How many synonyms are sufficient? And from which Corpus?** Previous results have demonstrated empirically that the proposed approach brings substantial improvements by just replacing the categorical name on the ground truth by a closely related synonym. Thus, we now want to evaluate the impact of the corpus selected, as well as the amount of synonyms needed to achieve satisfactory results. In particular, we evaluate the performance of the proposed approach (Fig 3), without adapters, i.e., 'A photo of [CLS\*].', when the optimal [CLS] token is selected from a set of potential synonyms extracted from different corpus: British National Corpus [11], Google News [29] and English Wikipedia. The first observation is that, while the use of different cor-

**Table 3.** Evaluation results on PASCAL VOC2012 *val* and *test* sets. The best results are in **bold**. Sup. denotes the weak supervision type.  $\mathcal{F}$  denotes full supervision.  $\mathcal{S}$  denotes saliency map supervision.  $\mathcal{I}$  denotes image-level supervision. Seg. denotes the segmentation network. Bac. denotes the backbone network for CAMs generation. V1: DeepLabV1. V2: DeepLabV2. V16: VGG-16 [54]. R50: ResNet-50 [20]. WR38: WideResNet38 [64]. Segmentation network pretrained with ImageNet otherwise using MS COCO dataset ( $\ddagger$ ). Approaches based on visual transformers (*last section*) and convolutional neural networks (*before-last section*) are separated, and best method in each is highlighted in **bold**.

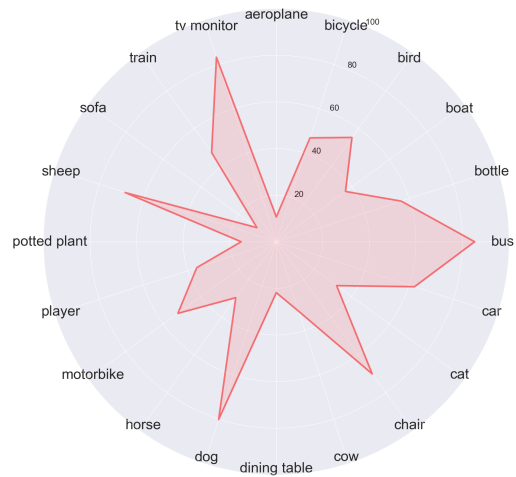
Sup.	Method	Seg.	Bac.	<i>val</i>	<i>test</i>
$\mathcal{F}$	DeepLabV2 <sub>TPAMI'18</sub> [7]	-	-	77.6	79.7
	WideResNet38 <sub>PR'19</sub> [65]	-	-	80.8	82.5
$\mathcal{I} + \mathcal{S}$	NSROM <sub>CVPR'21</sub> [70]	V2 $\ddagger$ -R101	V16	68.3	68.5
	DRS <sub>AAAI'21</sub> [27]	V2 $\ddagger$	V16	70.4	70.7
	EPS <sub>CVPR'21</sub> [35]	V2 $\ddagger$ -R101	WR38	70.9	70.8
	EDAM <sub>CVPR'21</sub> [63]	V2 $\ddagger$ -R101	WR38	70.9	70.6
	AuxSegNet <sub>ICCV'21</sub> [68]	WR38	-	69.0	68.6
	SANCE <sub>CVPR'22</sub> [68]	V2-R101	R50	72.0	72.9
$\mathcal{I}$	SEAM <sub>CVPR'20</sub> [60]	V3-R38	WR38	64.5	65.7
	BES <sub>ECCV'20</sub> [8]	V2-R101	R50	65.7	66.6
	SC-CAM <sub>CVPR'20</sub> [6]	V2-R101	WR38	66.1	65.9
	A <sup>2</sup> GNN <sub>TPAMI'21</sub> [72]	V2-R101	WR38	66.8	67.4
	VWE <sub>IJCAI'21</sub> [51]	V2-R101	R50	67.2	67.3
	AdvCAM <sub>CVPR'21</sub> [32]	V2-R101	R50	68.1	68.0
	VWL <sub>IJCV'22</sub> [52]	V2 $\ddagger$ -R101	R101	70.6	70.7
	SIPE <sub>CVPR'22</sub> [9]	V2-R38	R50	68.2	69.5
	CLIMS <sub>CVPR'22</sub> [66]	V2 $\ddagger$ -R101	R50	70.4	70.0
	AdvCAM+W-OoD <sub>CVPR'22</sub> [33]	V2-R101	R50	69.8	69.9
	SIPE <sub>CVPR'22</sub> [9]	V2-R101	R50	68.8	69.7
	RECAM <sub>CVPR'22</sub> [10]	V2-R101	R50	68.5	68.4
	AMN <sub>CVPR'22</sub> [34]	V2 $\ddagger$ -R101	R50	70.7	70.6
	Spatial-BCE <sub>ECCV'22</sub> [62]	V2-R101	R38	68.5	69.7
	ESOL <sub>NeurIPS'22</sub> [38]	V2-R101	R50	69.9	69.3
	<b>POLE (ours)</b>	V2 $\ddagger$ -R101	R50	<b>71.5</b>	<b>71.4</b>
$\mathcal{I}$	AFA <sub>CVPR'22</sub> [53]	MiT-B1	-	66.0	66.3
	MCTformer <sub>CVPR'22</sub> [69]	V1-R38	DeiT-S	<b>71.9</b>	<b>71.6</b>
	ViT-PCM <sub>ECCV'22</sub> [50]	V2-R101	ViT-B/16	70.3	70.9

pus increase the performance over the baseline ('A photo of [CLS].' in Table 1), synonyms from ChatGPT yield the best performance, regardless on the number of names requested. This may be explained by the larger and richer body of text, from a variety of sources, used to train ChatGPT. Next, we can observe that the quality of the generated CAMs typically augments with the number of synonyms (e.g., with ChatGPT we obtain a mIoU of 72.2 vs 73.6, from 2 and 4 synonyms, respectively). These results showcase how the most semantically related words, from a natural language standpoint, do not always yield the best performance. Indeed, as the performance increases with the number of synonyms included in our method, one can easily deduce that the synonyms newly added (less correlated than the first ones) may provide better supervisory signals for certain images. Additionally, we investigate the frequency that the actual ground truth category is selected as the [CLS] token, whose results are depicted in Fig. 4. The findings from this radar plots reveal that, surprisingly, the ground truth associated with most instances does not correspond to the most correlated category, which may explain the performance gains observed in our approach.

Further exploration on the choice of the synonyms across corpus can be found in Supplemental Material.

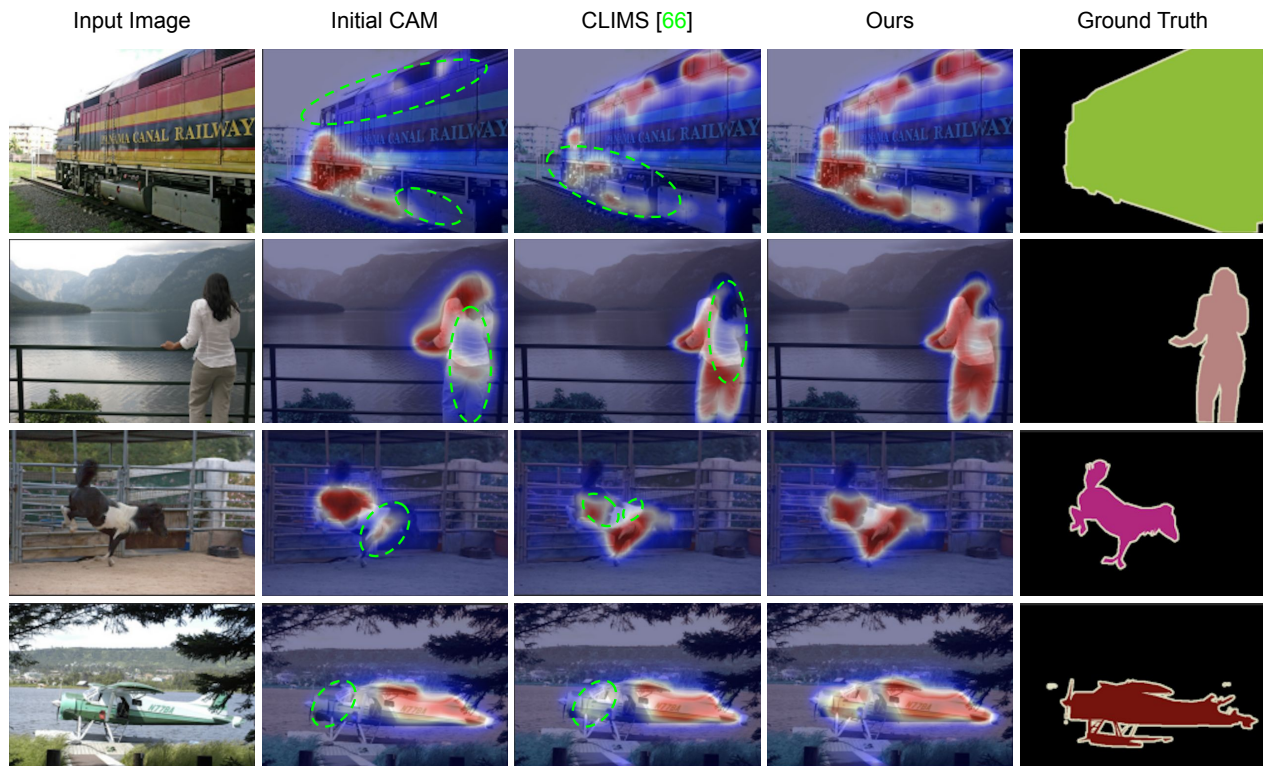


**Figure 3. Impact of the Corpus choice and number of synonyms selected.** ChatGPT offers the richer variety of synonyms, yielding the best results across other corpus. Furthermore, increasing the number of synonyms (from 2 up to 4) further improves the results. Note that the number of synonyms includes the categorical name from the ground truth and the requested close synonyms.



**Figure 4. What does CLIP think about the best [CLS]? Is the ground truth category chosen everytime? How likely is it that CLIP will select something different?** The plot summarises the percentage of cases where the ground truth category was chosen for an instance of that class. Thus, an inward point on the radial plot indicates that the number of instances where the ground truth category was chosen as the best [CLS] token is considerably low.

**On the impact of the different components.** We observed in Table 1 that the proposed yet simple strategy to optimize the category name achieves better performance than arguably more complex techniques that attempt to optimize the whole text prompt. In this section, we empirically motivate the use of the proposed adapters, as well as the choice of adding learnable parameters to control the importance of each term in Eq. 7, unlike the fixed hyperparameter used in the existing literature. In particular, table 4 reports these results, where the first observation is that adding the proposed adaptors results in slight improvement compared to



**Figure 5. Qualitative results of the initial class activation maps.** Green dotted lines ellipses are used to indicate missed regions by previous approaches (original CAMs and CLIMS [66]) compared to the proposed method. No refinement on the obtained CAMs is done (e.g., RW) to better illustrate the impact of our approach.

the model without them. In contrast, replacing the fixed vectors by learnable ones, the performance is further improved by 0.4%. Thus, the negligible increase in model complexity due to the adapters, and the performance gain observed, support the choices behind our approach.

**Table 4. Ablation on the main components.** Empirical results that validate the different components of the proposed methodology. A) Image label as [CLS], i.e., CLIMS [66]; B) Optimal [CLS] selected; C) Fixed adaptor; D) Learnable adaptor. Results are performed on the *train* set of PASCAL VOC2012.

A	B	C	D	mIoU (%)
✓				70.5
	✓			73.6 <sub>(+3.1)</sub>
	✓	✓		73.8 <sub>(+3.3)</sub>
	✓		✓	74.2 <sub>(+3.7)</sub>

**Qualitative results** of the obtained CAMs, compared to standard CAM generation and the related CLIMS [66] are depicted in Figure 5. We can observe that despite CLIMS somehow alleviates the under-segmentation problem in conventional CAMs, it still fails to cover larger target regions (see for example first and second columns). In contrast, POLE typically identifies better larger semantic regions re-

lated to the target class, resulting in more complete CAMs compared to related approaches.

## 5. Conclusions

In this work, we have investigated the potential of prompt tuning, an emerging strategy to adapt large pre-trained language-vision models, in the challenging task of weakly supervised semantic segmentation. Our empirical observations have demonstrated that simply replacing the text-token associated with the category name yields better segmentation performance than more complex prompt learning strategies focusing on optimizing the context, which dominate the literature in adapting models. More interestingly, we have observed that employing the corresponding image-level ground truth does not always lead to the best segmentation performance, and closely-related synonyms can indeed result in further performance gains. In light of these findings, we have introduced a simple yet efficient approach, POLE, that selects the most correlated class for a given image in order to generate a better text prompt. Comprehensive experiments have shown that the proposed approach can generate high quality pseudo-labels for WSSS, and achieve state-of-the-art performance in a popular WSSS benchmark.



## References

- [1] Information provided by chatgpt, an artificial intelligence language model developed by openai. Accessed on [March 23rd, 2023]. Generated text should be treated as such. [4](#)
- [2] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2209–2218, 2019. [5](#), [6](#)
- [3] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4981–4990, 2018. [2](#), [3](#), [6](#)
- [4] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4981–4990, 2018. [5](#), [6](#)
- [5] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *Proceedings of the European Conference on Computer Vision*, pages 549–565. Springer, 2016. [2](#)
- [6] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8991–9000, 2020. [3](#), [6](#), [7](#)
- [7] 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. *Transactions on Pattern Analysis and Machine Intelligence*, 40:834–848, 2018. [7](#), [13](#)
- [8] Liyi Chen, Weiwei Wu, Chenchen Fu, Xiao Han, and Yuntao Zhang. Weakly supervised semantic segmentation with boundary exploration. In *Proceedings of the European Conference on Computer Vision*, pages 347–362. Springer, 2020. [7](#)
- [9] Qi Chen, Lingxiao Yang, Jian-Huang Lai, and Xiaohua Xie. Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4288–4298, 2022. [6](#), [7](#)
- [10] Zhaozheng Chen, Tan Wang, Xiongwei Wu, Xian-Sheng Hua, Hanwang Zhang, and Qianru Sun. Class re-activation maps for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 969–978, 2022. [6](#), [7](#)
- [11] BNC Consortium et al. British national corpus. *Oxford Text Archive Core Collection*, 2007. [6](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. [3](#)
- [13] Yu Du, Fangyun Wei, Zihe Zhang, Miaojing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14084–14093, 2022. [3](#)
- [14] 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*, 88:303–308, 2009. [5](#)
- [15] Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4283–4292, 2020. [2](#)
- [16] Junsong Fan, Zhaoxiang Zhang, Tieniu Tan, Chunfeng Song, and Jun Xiao. CIAN: Cross-image affinity net for weakly supervised semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10762–10769, 2020. [3](#)
- [17] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. [5](#)
- [18] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2(3):4, 2021. [3](#)
- [19] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 991–998. IEEE, 2011. [5](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. [5](#), [7](#)
- [21] Qibin Hou, PengTao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. *Advances in Neural Information Processing Systems*, 31, 2018. [2](#), [3](#)
- [22] Shell Xu Hu, Da Li, Jan Stühmer, Minyoung Kim, and Timothy M Hospedales. Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9068–9077, 2022. [3](#)
- [23] Peng-Tao Jiang, Yuqi Yang, Qibin Hou, and Yunchao Wei. L2G: A simple local-to-global knowledge transfer framework for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16886–16896, 2022. [3](#)
- [24] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *Proceedings of the European Conference on Computer Vision*, pages 105–124. Springer, 2022. [2](#)
- [25] Hoel Kervadec, Jose Dolz, Meng Tang, Eric Granger, Yuri Boykov, and Ismail Ben Ayed. Constrained-CNN losses for weakly supervised segmentation. *Medical image analysis*, 54:88–99, 2019. [2](#)

- [26] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 876–885, 2017. 2
- [27] Beomyoung Kim, Sangeun Han, and Junmo Kim. Discriminative region suppression for weakly-supervised semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1754–1761, 2021. 7
- [28] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 695–711. Springer, 2016. 2
- [29] Andrei Kutuzov, Murhaf Fares, Stephan Oepen, and Erik Velldal. Word vectors, reuse, and replicability: Towards a community repository of large-text resources. In *Proceedings of the 58th Conference on Simulation and Modelling*, pages 271–276. Linköping University Electronic Press, 2017. 6
- [30] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6994–7003, 2021. 2, 3
- [31] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5267–5276, 2019. 2, 3
- [32] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4071–4080, 2021. 5, 6, 7
- [33] Jungbeom Lee, Seong Joon Oh, Sangdoon Yun, Junsuk Choe, Eunji Kim, and Sungroh Yoon. Weakly supervised semantic segmentation using out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16897–16906, 2022. 6, 7
- [34] Minhyun Lee, Dongseob Kim, and Hyunjung Shim. Threshold matters in wsss: manipulating the activation for the robust and accurate segmentation model against thresholds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4330–4339, 2022. 7
- [35] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5495–5505, 2021. 3, 7
- [36] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, 2021. 3
- [37] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *International Conference on Learning Representations*, 2021. 3
- [38] Jinlong Li, Zequn Jie, Xu Wang, Xiaolin Wei, and Lin Ma. Expansion and shrinkage of localization for weakly-supervised semantic segmentation. In *Advances in Neural Information Processing Systems*, 2022. 6, 7
- [39] Ruiwen Li, Zheda Mai, Chiheb Trabelsi, Zhibo Zhang, Jongseong Jang, and Scott Sanner. Transcam: Transformer attention-based cam refinement for weakly supervised semantic segmentation. *arXiv preprint arXiv:2203.07239*, 2022. 3
- [40] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL*, pages 4582–4597, 2021. 3
- [41] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3159–3167, 2016. 2
- [42] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35, 2023. 3
- [43] Robert Logan IV, Ivana Balažević, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. Cutting down on prompts and parameters: Simple few-shot learning with language models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2824–2835, 2022. 6
- [44] Nihal V Nayak, Peilin Yu, and Stephen H Bach. Learning to compose soft prompts for compositional zero-shot learning. In *International Conference on Learning Representations*, 2023. 3
- [45] Gaurav Patel and Jose Dolz. Weakly supervised segmentation with cross-modality equivariant constraints. *Medical Image Analysis*, 77:102374, 2022. 3
- [46] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1796–1804, 2015. 2
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 3, 4, 13
- [48] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 3
- [49] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. 3, 5
- [50] Simone Rossetti, Damiano Zappia, Marta Sanzari, Marco Schaerf, and Fiora Pirri. Max pooling with vision transform-

- ers reconciles class and shape in weakly supervised semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 446–463. Springer, 2022. 6, 7
- [51] Lixiang Ru, Bo Du, and Chen Wu. Learning visual words for weakly-supervised semantic segmentation. In *IJCAI*, volume 5, page 6, 2021. 7
- [52] Lixiang Ru, Bo Du, Yibing Zhan, and Chen Wu. Weakly-supervised semantic segmentation with visual words learning and hybrid pooling. *International Journal of Computer Vision*, 130(4):1127–1144, 2022. 6, 7
- [53] Lixiang Ru, Yibing Zhan, Baosheng Yu, and Bo Du. Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16846–16855, 2022. 6, 7
- [54] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 7
- [55] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3136–3145, 2019. 2
- [56] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 347–365. Springer, 2020. 3
- [57] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 507–522. Springer, 2018. 2
- [58] Feng Wang, Manling Li, Xudong Lin, Hairong Lv, Alexander G Schwing, and Heng Ji. Learning to decompose visual features with latent textual prompts. *arXiv preprint arXiv:2210.04287*, 2022. 3, 5, 6, 13
- [59] Xiang Wang, Sifei Liu, Huimin Ma, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation by iterative affinity learning. *International Journal of Computer Vision*, 128(6):1736–1749, 2020. 3
- [60] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020. 3, 6, 7
- [61] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1568–1576, 2017. 3
- [62] Tong Wu, Guangyu Gao, Junshi Huang, Xiaolin Wei, Xiaoming Wei, and Chi Harold Liu. Adaptive spatial-bce loss for weakly supervised semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 199–216. Springer, 2022. 7
- [63] Tong Wu, Junshi Huang, Guangyu Gao, Xiaoming Wei, Xiaolin Wei, Xuan Luo, and Chi Harold Liu. Embedded discriminative attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16765–16774, 2021. 3, 7
- [64] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019. 7
- [65] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019. 7
- [66] Jinheng Xie, Xianxu Hou, Kai Ye, and Linlin Shen. CLIMS: Cross language image matching for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4483–4492, 2022. 2, 3, 4, 5, 6, 7, 8, 13, 14
- [67] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 3
- [68] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6984–6993, 2021. 7
- [69] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Multi-class token transformer for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4310–4319, 2022. 3, 6, 7, 13
- [70] Yazhou Yao, Tao Chen, Guo-Sen Xie, Chuanyi Zhang, Fumin Shen, Qi Wu, Zhenmin Tang, and Jian Zhang. Non-salient region object mining for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2623–2632, 2021. 7
- [71] Sung-Hoon Yoon, Hyeokjun Kweon, Jegyeong Cho, Shinjeong Kim, and Kuk-Jin Yoon. Adversarial erasing framework via triplet with gated pyramid pooling layer for weakly supervised semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 326–344. Springer, 2022. 6
- [72] Bingfeng Zhang, Jimin Xiao, Jianbo Jiao, Yunchao Wei, and Yao Zhao. Affinity attention graph neural network for weakly supervised semantic segmentation. *Transactions on Pattern Analysis and Machine Intelligence*, 2021. 7
- [73] Fei Zhang, Chaochen Gu, Chenyue Zhang, and Yuchao Dai. Complementary patch for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7242–7251, 2021. 3
- [74] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *Proceedings of the European Conference on Computer Vision*, pages 493–510. Springer, 2022. 5
- [75] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE/CVF Con-*

- ference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016. [2](#)
- [76] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [13](#)
- [77] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [2](#), [3](#), [5](#), [13](#)
- [78] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11175–11185, 2023. [3](#)