

# Token Contrast for Weakly-Supervised Semantic Segmentation

Lixiang Ru<sup>1</sup> Heliang Zheng<sup>2</sup> Yibing Zhan<sup>2</sup> Bo Du<sup>1\*</sup>

<sup>1</sup> Institute of Artificial Intelligence, School of Computer Science,  
National Engineering Research Center for Multimedia Software and Hubei Key Laboratory  
of Multimedia and Network Communication Engineering, Wuhan University, China.

<sup>2</sup> JD Explore Academy, China

{rulixiang, dubo}@whu.edu.cn {zhengheliang, zhanyibing}@jd.com

## Abstract

*Weakly-Supervised Semantic Segmentation (WSSS) using image-level labels typically utilizes Class Activation Map (CAM) to generate the pseudo labels. Limited by the local structure perception of CNN, CAM usually cannot identify the integral object regions. Though the recent Vision Transformer (ViT) can remedy this flaw, we observe it also brings the over-smoothing issue, i.e., the final patch tokens incline to be uniform. In this work, we propose Token Contrast (ToCo) to address this issue and further explore the virtue of ViT for WSSS. Firstly, motivated by the observation that intermediate layers in ViT can still retain semantic diversity, we designed a Patch Token Contrast module (PTC). PTC supervises the final patch tokens with the pseudo token relations derived from intermediate layers, allowing them to align the semantic regions and thus yield more accurate CAM. Secondly, to further differentiate the low-confidence regions in CAM, we devised a Class Token Contrast module (CTC) inspired by the fact that class tokens in ViT can capture high-level semantics. CTC facilitates the representation consistency between uncertain local regions and global objects by contrasting their class tokens. Experiments on the PASCAL VOC and MS COCO datasets show the proposed ToCo can remarkably surpass other single-stage competitors and achieve comparable performance with state-of-the-art multi-stage methods. Code is available at <https://github.com/rulixiang/ToCo>.*

## 1. Introduction

To reduce the expensive annotation costs of deep semantic segmentation models, weakly-supervised semantic segmentation (WSSS) is proposed to predict pixel-level predictions with only weak and cheap annotations, such as image-level labels [2], points [4], scribbles [51] and bounding

\*Corresponding author. This work was done when Lixiang Ru was a research intern at JD Explore Academy.

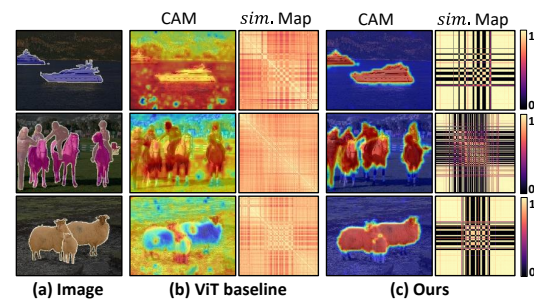


Figure 1. **The generated CAM and the pairwise cosine similarity of patch tokens (*sim. map*).** Our method can address the over-smoothing issue well and produce accurate CAM. Here we use ViT-Base.

boxes [23]. Among all these annotation forms, the image-level label is the cheapest and contains the least information. This work also falls in the field of WSSS using only image-level labels.

Prevalent works of WSSS using image-level labels typically derive Class Activation Map (CAM) [53] or its variants [35] as pseudo labels. The pseudo labels are then processed with alternative refinement methods [1, 2] and used to train regular semantic segmentation models. However, CAM is usually flawed since it typically only identifies the most discriminative semantic regions, severely weakening the final performance of semantic segmentation [1, 19, 43]. The recent works [15, 34, 47] show one reason is that previous methods usually generate CAM with CNN, in which convolution only perceives local features and fails to activate the integral object regions. To ameliorate this problem and generate more accurate pseudo labels for WSSS, these works propose solutions based on the recent Vision Transformer (ViT) architecture [12], which inherently models the global feature interactions with self-attention blocks.

However, as demonstrated in [29, 42], self-attention in ViT is essentially a low-pass filter, which inclines to reduce the variance of input signals. Therefore, stacking self-attention blocks is equivalent to repeatedly performing spatial smoothing operations, which encourages the patch to-

kens in ViT to be uniform [16, 36], *i.e.*, over-smoothing. We observe that the over-smoothing issue particularly impairs the WSSS task, since CAM used to derive pseudo labels relies on the output features (*i.e.* patch tokens). As shown in Figure 1, due to over-smoothing, the pairwise cosine similarities of the patch tokens are close to 1, suggesting the learned representations of different patch tokens are almost uniform. The generated CAM thus tends to assign different image regions with the monotonous semantic label. Though several recent works have explored the ViT architecture for WSSS [34, 39, 47], they typically overlook the over-smoothing issue of patch tokens, leaving this problem unresolved.

In this work, we empirically observe that ViT smooths the patch tokens progressively, *i.e.* the learned representations in intermediate layers can still preserve the semantic diversity. Therefore, we propose a Patch Token Contrast (PTC) module to address the over-smoothing issue by supervising the final patch tokens with intermediate layer knowledge. Specifically, in the PTC module, we simply add an additional classifier in an intermediate layer to extract the auxiliary CAM and the corresponding pseudo pairwise token relations. By supervising the pairwise cosine similarities of final patch tokens with the pseudo relations, PTC can finely counter the over-smoothing issue and thus produce high-fidelity CAM. As shown in Figure 1, our method can generate CAM that aligns well with the semantic object regions. The pairwise cosine similarities also coincide with the corresponding semantics. In addition, to further differentiate the uncertain regions in generated CAM, inspired by the property that the class token in ViT can inherently aggregate high-level semantics [6, 15], we also propose a Class Token Contrast (CTC) module. In CTC, we first randomly crop local images from uncertain regions (background regions), and minimize (maximize) the representation difference between the class tokens of local and global images. As a result, CTC can facilitate the local-to-global representation consistency of semantic objects and the discrepancy between foreground and background, benefiting the integral and accurate object activation in CAM. Finally, based on the proposed PTC and CTC, we build Token Contrast (ToCo) for WSSS and extend it to the single-stage WSSS framework [34].

Overall, our contributions in this work include the following aspects.

- We propose Patch Token Contrast (PTC) to address the over-smoothing issue in ViT. By supervising the final tokens with intermediate knowledge, PTC can counter the patch uniformity and significantly promote the quality of pseudo labels for WSSS.
- We propose Class Token Contrast (CTC), which contrasts the representation of global foregrounds and local uncertain regions (background) and facilitates the object activation completeness in CAM.

ation completeness in CAM.

- The experiments on the PASCAL VOC [14] and MS COCO dataset [26] show that the proposed ToCo can significantly outperform SOTA single-stage WSSS methods and achieve comparable performance with multi-stage competitors.

## 2. Related Work

**Weakly-Supervised Semantic Segmentation.** Weakly-Supervised Semantic Segmentation (WSSS) using image-level labels typically generates CAM as the initial pseudo labels. A typical drawback of CAM is that it usually only activates the most discriminative regions. To address this drawback, recent works proposed various training schemes, such as erasing [44], online attention accumulation [18] and cross-image semantic mining [38]. [7, 33, 43] propose to leverage auxiliary tasks to regularize the training objective, such as visual words learning [33], subcategory exploration [7], and scale in-variance regularization [43]. [24, 49] utilize extra saliency maps as supervision to suppress the background regions and mine the non-salient objects. [22, 37, 46] counter the problem of semantic co-occurrence via distilling knowledge from CLIP [30], decoupling object context [37] and comparing the out-of-distribution images [22], respectively. [9, 13, 54] contrast the pixel and prototype representations to encourage the integral activation of the object regions. Typically, these methods are built upon the CNN network, inheriting the locality flaw. In this work, we explore ViT for WSSS to avoid this drawback and achieve integral object activation.

**Vision Transformer for WSSS.** Vision Transformer (ViT) has achieved great success in various vision tasks [5, 12, 20, 45, 48]. Some recent works also introduce ViT to WSSS [15, 34, 39, 47]. Inspired by the property that the class token in ViT can capture the foreground information [6], TS-CAM [15] extracts the class-agnostic attention map and couples it with the naive semantic-aware CAM. [39] proposes to derive the gradients of the attention map and extracts the attention of class token w.r.t. other tokens as the class-specific maps. MCTformer [47] embeds multiple class tokens and enforces them learning the activation maps of different classes. AFA [34] proposes to learn reliable semantic affinity from the attention blocks to refine the initial coarse labels. However, these methods typically overlook the over-smoothing issue of ViT. Besides, they need to modify the ViT architecture or extract the costly gradients at the inference stage. In this work, we propose ToCo to solve the over-smoothing issue without modification to architecture and further unlock the potential of ViT on the WSSS task.

## 3. Preliminaries

In this section, we briefly introduce the preliminary knowledge of class activation map (CAM), Vision Trans-




Figure 2. **The average pairwise cosine similarity of patch tokens in each Transformer block.** The cosine similarity is computed on the VOC `train` set. Here we use the ViT-Base (ViT-B) [12] architecture which includes 12 Transformer blocks.

former (ViT), and over-smoothing in ViT.

### 3.1. Class Activation Map

CAM [53] is proposed to identify the activated regions when a classification network predicts an image. Due to its efficacy and simplicity, CAM has been widely used to generate the initial pseudo labels for WSSS. Specifically, given an image, its feature maps  $\mathbf{F} \in \mathbb{R}^{hw \times d}$  are extracted with a classification network (CNN or ViT), where  $hw$  and  $d$  denote the number of spatial and channel dimension, respectively. The CAM is computed by weighting and summing the feature maps with the weights  $\mathbf{W} \in \mathbb{R}^{c \times d}$  in the classification layer, where  $c$  is the number of semantic classes. The `relu` function and `max` normalization are then applied to eliminate the negative activation and scale the CAM to  $[0, 1]$ . Therefore, CAM for class  $c$  is thus calculated as

$$\text{CAM}_c(\mathbf{F}, \mathbf{W}) = \frac{\text{relu}(\mathbf{M}_c)}{\max(\text{relu}(\mathbf{M}_c))}, \mathbf{M}_c = \sum_i \mathbf{W}_{c,i} \mathbf{F}_{:,i}. \quad (1)$$

A background threshold  $\beta$  is usually used to differentiate the background and foreground regions.

### 3.2. Vision Transformer & Over-smoothing

An ViT firstly splits an image into patches to form the initial patch tokens. Then the patch tokens are concatenated with an extra learnable class token and feed into the Transformer encoder to obtain the final patch and class tokens. As the key component, in each Transformer block, the multi-head self-attention (MHSA) is used to perform the global feature interaction. However, due to the low-pass property of self-attention [29, 42], after multiple Transformer blocks, the output patch tokens incline to be uniform, which severely affects the CAM according to Equation 1.

In Figure 2, we visualized the pairwise cosine similarity of the patch tokens generated in each Transformer block. Figure 2 shows that the patch tokens in the late layers are highly similar, while the early layers can still preserve the

semantic diversity. This observation motivates us to address the over-smoothing issue by supervising the final layer tokens with knowledge from intermediate layers.

## 4. Methodology

This section elaborates on the proposed method, *i.e.* Token Contrast (ToCo) for WSSS. We first introduce the overall framework of ToCo. Then the Patch Token Contrast (PTC) and Class Token Contrast (CTC) are proposed to address the over-smoothing issue and further exploit the virtue of ViT for WSSS, respectively. Finally, we present the training objective of ToCo and how to plug it into the single-stage WSSS framework.

### 4.1. Overview

As illustrated in Figure 3, ToCo uses an auxiliary classification layer in the ViT encoder to produce the auxiliary CAM. The auxiliary CAM is subsequently leveraged to generate the auxiliary pseudo labels and guide the PTC module. Meanwhile, it’s also used to produce proposals to crop positive and negative local images for the CTC module. The final CAM is obtained with a classification layer and used to generate the final pseudo labels.

### 4.2. Patch Token Contrast

The objective of the Patch Token Contrast (PTC) module is to address the over-smoothing issue of the final patch tokens. As aforementioned, since the intermediate layers can still preserve the semantic diversity of patch tokens, in PTC, we leverage knowledge from the intermediate layer, *i.e.* reliable pairwise token relations in Figure 3, to supervise the final patch tokens.

Specifically, an input image  $\mathbf{X}$  is firstly tokenized to construct initial patch tokens and then passed through the Transformer encoder. For a given intermediate layer, the output patch tokens are denoted as  $\mathbf{F}^m \in \mathbb{R}^{n \times d}$ , where  $n$  and  $d$  denotes the number of tokens and the feature dimension, respectively. As shown in Figure 3, to extract the semantic-aware knowledge, we add an auxiliary classification layer to perform classification and generate CAM. In practice, we observe that not all intermediate layers produce satisfactory CAM for supervising PTC, since late layers incline to smooth the patch tokens while early layers may fail to capture high-level semantics. The choice of the intermediate layer will be discussed in Section 5.3.

In the auxiliary classification head, the patch tokens  $\mathbf{F}^m$  are firstly aggregated via global max-pooling (GMP) as suggested in [34] and then projected with a fully-connected layer parameterized with  $\theta^m$  to calculate the auxiliary classification loss  $\mathcal{L}_{cls}^m$ . Therefore, the auxiliary CAM is computed as

$$\mathbf{M}^m = \text{CAM}(\mathbf{F}^m, \theta^m). \quad (2)$$

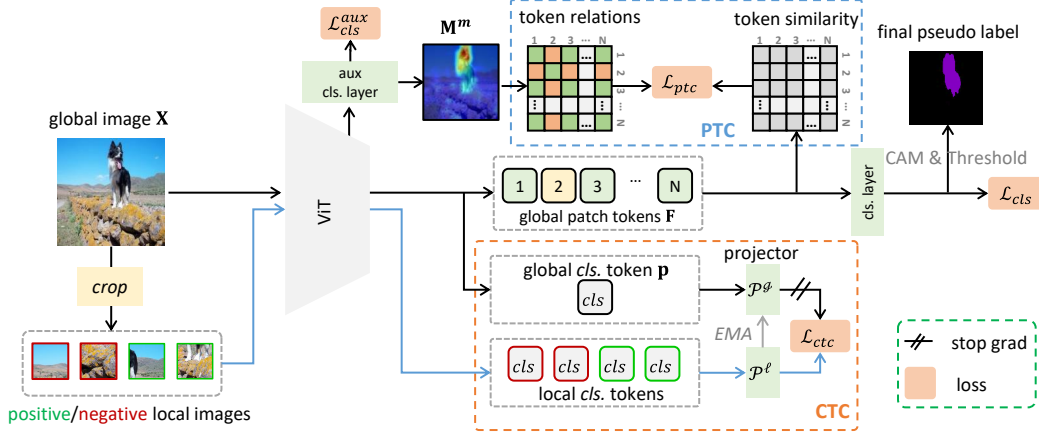


Figure 3. **The overall framework of ToCo.** ToCo firstly uses an additional classification layer (*cls.* layer) to produce the auxiliary CAM (*aux.* CAM). In the PTC module, the *aux.* CAM is used to derive token relations and supervise the pairwise token similarities of final patch tokens to address the over-smoothing issue. In the CTC module, the class tokens of the negative/positive images will be projected and contrasted with the global class token to further differentiate the low-confidence regions in CAM. The pseudo label is generated with the final CAM.

We then use two background thresholds  $\beta_l, \beta_h$  ( $0 < \beta_l < \beta_h < 1$ ) to segment  $\mathbf{M}^m$  to the pseudo token label  $\mathbf{Y}^m$ , which consists of reliable foreground, background and uncertain regions.

**Patch Token Contrast Loss.** The generated token label  $\mathbf{Y}^m$  is used to derive reliable pairwise relations for supervising the final patch tokens. Specifically, if two tokens share the same semantic label, they are labeled as positive pairs; otherwise, they are labeled as negative pairs. In addition, to ensure reliability, we only consider two tokens that both belong to the reliable foreground or background regions and ignore the uncertain regions. To remedy the over-smoothing issue, we maximize the similarity of two final patch tokens that belong to positive pairs and minimize the similarity otherwise. Let  $\mathbf{F} \in \mathbb{R}^{n \times d}$  be the final layer patch tokens, the loss function for the PTC module is then constructed as

$$\begin{aligned} \mathcal{L}_{ptc} = & \frac{1}{N^+} \sum_{\mathbf{Y}_i = \mathbf{Y}_j} (1 - \text{CosSim}(\mathbf{F}_i, \mathbf{F}_j)) \\ & + \frac{1}{N^-} \sum_{\mathbf{Y}_i \neq \mathbf{Y}_j} \text{CosSim}(\mathbf{F}_i, \mathbf{F}_j), \end{aligned} \quad (3)$$

where  $\text{CosSim}(\cdot, \cdot)$  computes the cosine similarity and  $N^+/N^-$  counts the number of positive/negative pairs. However, minimizing the original cosine similarity cannot ensure the diversity [10, 16], since a token pair with minus cosine similarity (e.g.,  $-1$ ) could be highly correlated. Therefore, in practice, in Equation 3, we use the absolute cosine similarity instead of the original form. By minimizing Equation 3, the representations of positive tokens are encouraged to be more consistent, while the negative tokens pairs are more discriminative, so that the over-smoothing issue can be well addressed.

### 4.3. Class Token Contrast

Addressing the over-smoothing with PTC can drive ViT to generate compelling CAM and pseudo labels. However, inevitably, there are still some less discriminative object regions that are hard to differentiate in CAM. Inspired by the property that class tokens in ViT can aggregate the high-level semantics [6, 15], we design a Class Token Contrast (CTC) module to facilitate the representation consistency between the local non-salient regions and the global object, which can further enforce more object regions to be activated in CAM.

As illustrated in Figure 4, given an image, we first randomly crop the local images from the uncertain regions specified by its auxiliary CAM. Since the class token in ViT captures the information of semantic objects [6, 15], the class tokens of global and local images aggregate the information of the global and the local objects, respectively. By minimizing the difference between global and local class tokens, the representation of entire object regions can be more consistent.

To counter the case that the cropped local images may contain few/no foreground objects, as shown in Figure 3, we also crop some local images from background regions. By maximizing the difference between class tokens of global image and local background regions, the foreground-background discrepancy can be also facilitated. In practice, we randomly crop a fixed number of local images and assign them as positive (from uncertain regions) or negative (from background regions) with the guidance of  $\mathbf{Y}^m$  in Section 4.2.

Specifically, the global and local class tokens are first passed through the projection head  $\mathcal{P}^g$  and  $\mathcal{P}^l$ , respectively, which consist of linear layers and an L2 normaliza-




Figure 4. Illustration of the crop method in Figure 3.

tion layer. Assuming  $\mathbf{p}$  denotes the projected global class token, and  $Q^+/Q^-$  denotes the set of projected local class tokens cropped from uncertain/background regions, the objective of CTC is to minimize/maximize the difference between  $\mathbf{p}$  and the local class tokens in  $Q^+/Q^-$ . Here we use the InfoNCE loss [27] as the objective, *i.e.*

$$\mathcal{L}_{ctc} = \frac{1}{N^+} \sum_{\mathbf{q}^+} \log \frac{e^{(\mathbf{p}^\top \mathbf{q}^+ / \tau)}}{e^{(\mathbf{p}^\top \mathbf{q}^+ / \tau)} + \sum_{\mathbf{q}^-} e^{(\mathbf{p}^\top \mathbf{q}^- / \tau)} + \epsilon}, \quad (4)$$

where  $\mathbf{q}^+ \in Q^+$ ,  $\mathbf{q}^- \in Q^-$ ,  $N^+$  counts the number of  $Q^+$ ,  $\tau$  is the temperature factor, and  $\epsilon$  is a small positive value for stability. It's noted that CTC aims to enforce the local view representation to align the global view's. Therefore, we stop the gradient of the projection head  $\mathcal{P}^g$ . To update  $\mathcal{P}^g$ , we use the exponential moving average (EMA), *i.e.*,  $\theta^g \leftarrow \rho \theta^g + (1 - \rho) \theta^l$ , where  $\rho$  is the momentum factor,  $\theta^g$  and  $\theta^l$  are parameters from  $\mathcal{P}^g$  and  $\mathcal{P}^l$ , respectively.

#### 4.4. ToCo for WSSS

**Training Objective.** As shown in Figure 3, except the token contrast losses  $\mathcal{L}_{ptc}$  and  $\mathcal{L}_{ctc}$ , the training objective of the proposed ToCo also includes the classification loss  $\mathcal{L}_{cls}$  and auxiliary classification loss  $\mathcal{L}_{cls}^m$ . Following the common practice, we use the multi-label soft margin loss for both  $\mathcal{L}_{cls}$  and  $\mathcal{L}_{cls}^m$ . The optimization objective of ToCo is the weighted sum of these loss terms:

$$\mathcal{L}_{toco} = \mathcal{L}_{cls} + \mathcal{L}_{cls}^m + \lambda_1 \mathcal{L}_{ptc} + \lambda_2 \mathcal{L}_{ctc}. \quad (5)$$

**Single-Stage WSSS.** We plug the proposed ToCo into the single-stage WSSS framework. Specifically, the pseudo labels produced by ToCo are then refined with a pixel-adaptive refinement module (PAR) [34] to align the low-level semantic boundaries. The refined pseudo labels will be used to supervise the segmentation decoder. We use the common cross-entropy loss as the segmentation loss  $\mathcal{L}_{seg}$ . The overall training objective should thus include  $\mathcal{L}_{seg}$ , *i.e.*  $\mathcal{L} = \mathcal{L}_{toco} + \lambda_3 \mathcal{L}_{seg}$ . Following previous single-stage WSSS works [28, 33], we also use an extra regularization loss term [40] to enforce the spatial consistency of the predicted segmentation masks.

Method	Backbone	train	val
RRM [50] <small>AAAI'2020</small>	WR38	–	65.4
1Stage [3] <small>CVPR'2020</small>	WR38	66.9	65.3
AA&LR [52] <small>ACM MM'2021</small>	WR38	68.2	65.8
SLRNet [28] <small>IJCV'2022</small>	WR38	67.1	66.2
AFA [34] <small>CVPR'2022</small>	MiT-B1	68.7	66.5
ViT-PCM [32] <small>ECCV'2022</small>	ViT-B <sup>†</sup>	67.7	66.0
ViT-PCM + CRF [32] <small>ECCV'2022</small>	ViT-B <sup>†</sup>	71.4	69.3
<b>ToCo</b>	ViT-B	72.2	70.5
<b>ToCo<sup>†</sup></b>	ViT-B <sup>†</sup>	<b>73.6</b>	<b>72.3</b>

Table 1. Evaluation of pseudo labels. <sup>†</sup> denotes using ImageNet-21k [31] pretrained parameters.

## 5. Experiments

### 5.1. Experimental Settings

**Datasets.** We evaluate the proposed method on the PASCAL VOC 2012 [14] and MS COCO 2014 dataset [26]. Following common practice, VOC 2012 dataset is further augmented with the SBD dataset [17]. The train, val, and test set of the augmented dataset consist of 10582, 1449, and 1456 images, respectively. For COCO 2014 dataset, the train and val set consist of about 82k and 40k images, respectively. In the training stage, we only use image-level labels. By default, we report mIoU as the evaluation metric.

**Network Architectures.** We use the ViT-base (ViT-B) [12] as the backbone, which is initialized with ImageNet pretrained weights [31]. To ensure the backbone accepts input images of arbitrary size, the pos\_embedding will be resized to input size via bilinear interpolation. The projection heads in the CTC module, *i.e.*  $\mathcal{P}^g$  and  $\mathcal{P}^l$  in Figure 3, consist of 3 linear layers and an L2-normalization layer. The parameters in projection heads are randomly initialized. We use a simple segmentation head as the decoder, which consists of two  $3 \times 3$  convolutional layers (with a dilation rate of 5) and a  $1 \times 1$  prediction layer.

**Implementation Details.** We train ToCo with an AdamW optimizer. The learning rate linearly increases to  $6e^{-5}$  in the first 1500 iterations and decays with a polynomial scheduler for later iterations. The warm-up and decay rates are set as  $1e^{-6}$  and 0.9, respectively. For experiments on the VOC dataset, the batch size and total iterations are set as 4 and 20k, respectively. The crop size of global and local view images are  $448^2$  and  $96^2$ , respectively. Besides, we follow the multi-crop and data augmentation strategy in [6] for global and local views. By default, the background thresholds ( $\beta_l, \beta_h$ ) are set as (0.25, 0.7). The temperature factors  $\tau$  in Equation 4 is 0.5. The momentum factor for the EMA process in the CTC module is set as 0.9. The weight factors ( $\lambda_1, \lambda_2, \lambda_3$ ) of the loss terms in Section 4.4 are set as (0.2, 0.5, 0.1). In the inference stage, following the common practice in semantic segmentation [8], we use multi-scale




Figure 5. **Visualization of CAM.** From left to right, the CAM is generated with CNN baseline, AFA [34], ViT baseline, ViT with PTC, ViT with PTC and CTC, and the proposed ToCo.

testing and dense CRF processing.

For the experiments on the COCO dataset, the network is trained for 80k iterations with a batch size of 8. The background thresholds  $(\beta_l, \beta_h)$  are set as (0.25, 0.65), while other settings remain the same. The impact of hyper-parameters will be presented in Section 5.3 and Supplementary Material.

## 5.2. Experimental Results

**Pseudo Labels.** We first visualize the generated CAM with ToCo in Figure 5. As shown in Figure 5, our method can remarkably produce more integral and accurate CAM than the CNN methods and the recent single-stage WSSS method, *i.e.* AFA [34]. Compared to the ViT baseline, our ToCo also remedies the over-smoothing issue well.

We report the quantitative evaluation results of the pseudo labels generated with CAM in Table 1. The results are evaluated on the `train` and `val` set of the VOC dataset and compared with recent WSSS methods. Since ViT is initially pretrained on ImageNet-21k, for a fair comparison with other methods, we also report the results of ToCo with the ImageNet-1k pretrained weights, *i.e.* DeiT [41]. Table 1 shows that ToCo can produce higher quality pseudo labels than the competitors with both the ImageNet-1k and ImageNet-21k pretrained weights. Particularly, ToCo can significantly outperform ViT-PCM [32] which also uses ViT-B<sup>†</sup>, even though the latter employs additional CRF

<sup>1</sup><http://host.robots.ox.ac.uk:8080/anonymous/KV9AQH.html>

<sup>2</sup><http://host.robots.ox.ac.uk:8080/anonymous/GNTBBZ.html>

	<i>Sup.</i>	<i>Net.</i>	VOC		COCO
			val	test	val
<b>Multi-stage WSSS methods.</b>					
RIB [21] <small>NeurIPS'2021</small>	$\mathcal{I} + \mathcal{S}$	DL-V2	70.2	70.0	–
EPS [24] <small>CVPR'2021</small>	$\mathcal{I} + \mathcal{S}$	DL-V2	71.0	71.8	–
L2G [19] <small>CVPR'2022</small>	$\mathcal{I} + \mathcal{S}$	DL-V2	72.1	71.7	44.2
RCA [54] <small>CVPR'2022</small>	$\mathcal{I} + \mathcal{S}$	DL-V2	72.2	72.8	36.8
Du <i>et al.</i> [13] <small>CVPR'2022</small>	$\mathcal{I} + \mathcal{S}$	DL-V2	72.6	73.6	–
RIB [21] <small>NeurIPS'2021</small>	$\mathcal{I}$	DL-V2	68.3	68.6	43.8
ReCAM [11] <small>CVPR'2022</small>	$\mathcal{I}$	DL-V2	68.4	68.2	45.0
VWL [33] <small>IJCV'2022</small>	$\mathcal{I}$	DL-V2	69.2	69.2	36.2
W-OoD [22] <small>CVPR'2022</small>	$\mathcal{I}$	WR38	70.7	70.1	–
MCTformer [47] <small>CVPR'2022</small>	$\mathcal{I}$	WR38	71.9	71.6	42.0
ESOL [25] <small>NeurIPS'2022</small>	$\mathcal{I}$	DL-V2	69.9	69.3	42.6
<b>Single-stage WSSS methods.</b>					
RRM [50] <small>AAAI'2020</small>	$\mathcal{I}$	WR38	62.6	62.9	–
1Stage [3] <small>CVPR'2020</small>	$\mathcal{I}$	WR38	62.7	64.3	–
AFA [33] <small>CVPR'2022</small>	$\mathcal{I}$	MiT-B1	66.0	66.3	38.9
SLRNet [28] <small>IJCV'2022</small>	$\mathcal{I}$	WR38	67.2	67.6	35.0
<b>ToCo</b>	$\mathcal{I}$	ViT-B	69.8	70.5 <sup>1</sup>	41.3
<b>ToCo<sup>†</sup></b>	$\mathcal{I}$	ViT-B <sup>†</sup>	<b>71.1</b>	<b>72.2<sup>2</sup></b>	<b>42.3</b>

Table 2. **Semantic Segmentation Results.** *Sup.* denotes the supervision type.  $\mathcal{I}$ : Image-level labels;  $\mathcal{S}$ : Saliency maps. *Net.* denotes the backbone network (for single-stage methods) and the semantic segmentation network (for multi-stage methods). <sup>†</sup> denotes using ImageNet-21k [31] pretrained parameters.

post-processing.

**Semantic Segmentation Results.** The semantic segmentation results on the VOC and COCO datasets are reported in Table 2. The proposed ToCo achieves 71.1%, 72.2%




Figure 6. Semantic segmentation results on the VOC and COCO dataset. The last column shows a failure case.

Method	Backbone	val ( $\mathcal{F}$ )	val ( $\mathcal{I}$ )	ratio
1Stgae [3]	WR38	80.8	62.7	77.59%
AFA [34]	MiT-B1	78.7	66.0	83.86%
SLRNet [28]	WR38	80.8	67.2	83.17%
ToCo	ViT-B	80.5	69.8	<b>86.71%</b>
	ViT-B <sup>†</sup>	<b>82.3</b>	<b>71.1</b>	86.39%

Table 3. The fully-supervised counterparts of single-stage WSSS methods.  $\mathcal{F}$  /  $\mathcal{I}$  denotes the pixel- / image-level supervision. <sup>†</sup> denotes using ImageNet-21k [31] pretrained parameters.

and 42.3% mIoU on the VOC val, test and COCO val set, respectively, which largely outperform previous single-stage methods. Particularly, our single-stage ToCo achieves higher mIoU than multi-stage WSSS methods using image-level labels only. Other multi-stage methods using extra saliency maps ( $\mathcal{I} + \mathcal{S}$ ) only slightly outperform ToCo.

In Figure 6, we visualize and compare the predicted segmentation masks of ToCo, AFA [34], and the ground truth labels. Figure 6 shows ToCo can produce more accurate segmentation masks than AFA. Our single-stage WSSS results are also very close to the ground truth.

### 5.3. Ablation and Analysis

**Ablation.** To investigate the impact of the proposed PTC and CTC, in Table 4, we report the performance of the generated final CAM ( $\mathbf{M}$ ), auxiliary CAM ( $\mathbf{M}^m$ ) and the semantic segmentation results ( $Seg.$ ). The results are evaluated on the VOC val set.

We first show that due to over-smoothing, training a baseline ViT with the classification loss (*i.e.*  $\mathcal{L}_{cls}$ ) cannot produce reasonable CAM. Besides, adding an auxiliary classifier (*i.e.*  $\mathcal{L}_{cls} + \mathcal{L}_{cls}^m$ ) can help derive the auxiliary CAM from the intermediate layer, but it still cannot address the over-smoothing issue. The proposed PTC module can finely address this issue and significantly improve the quality of the generated CAM  $\mathbf{M}$ , *i.e.* from 27.9% mIoU to 62.5% mIoU. Table 4 also shows that the improvements in the final layer also benefit the intermediate CAM  $\mathbf{M}^m$ ,

Method	PTC	CTC	$\mathcal{L}_{seg}$	$\mathcal{L}_{reg}$	$\mathbf{M}$	$\mathbf{M}^m$	$Seg.$
$\mathcal{L}_{cls}$					27.8	–	–
$\mathcal{L}_{cls} + \mathcal{L}_{cls}^m$					27.9	53.8	–
$\mathcal{L}_{cls} + \mathcal{L}_{cls}^m$	✓				62.5	57.8	–
	✓	✓			67.2	60.7	–
	✓	✓	✓		69.9	61.2	66.6
	✓	✓	✓	✓	<b>70.5</b>	<b>62.5</b>	<b>68.1</b>

Table 4. Ablation Study.  $\mathbf{M}^m$ : auxiliary CAM from the intermediate layer;  $\mathbf{M}$ : CAM from the final layer;  $Seg.$ : semantic segmentation results.

improving  $\mathbf{M}^m$  to 57.8% mIoU. The CTC module further improves the quality of pseudo labels by 4.7% mIoU. It’s noted that the segmentation loss also improves the CAM’s quality, since we use PAR [33] to refine the pseudo labels which enforce the pseudo labels to align better with the low-level object boundaries. The regularization term  $\mathcal{L}_{reg}$  also brings improvements to the semantic segmentation results.

**Analysis of PTC.** To demonstrate that PTC addresses the over-smoothing issue well, in Figure 2, we compare the average pairwise cosine similarity of patch tokens of ViT with and without PTC. We show that in the late layers, the average cosine similarity with PTC is remarkably lower than the baseline. Note that the cosine similarity values decrease after the 10<sup>th</sup> layer, since we choose the 10<sup>th</sup> layer to produce the auxiliary labels and supervise the final CAM. We also visualize the generated CAM with and without PTC in Figure 5. Figure 5 shows PTC helps produce reasonable CAM, which coincides with Table 4.

**Analysis of CTC.** The motivation of CTC is to encourage the local-to-global consistency of semantic objects. In Figure 7, we visualize the attention map of class token w.r.t. other patch tokens in ToCo. We show that for both the global and local view, the class token finely captures the foreground object information. Moreover, the class token of the local view can learn the less salient regions that are usually ignored in the global view, which fulfils our intention. In Figure 5, we also visualize the generated CAM with

Block	M	M <sup>m</sup>	Seg.
#8	63.5	48.1	60.5
#9	67.6	55.1	64.8
#10	70.5	62.5	68.1
#11	43.1	45.2	40.3

(a) Index of auxiliary block.

Size	M	Seg.
64 <sup>2</sup>	68.1	65.5
80 <sup>2</sup>	69.3	67.3
96 <sup>2</sup>	70.5	68.1
112 <sup>2</sup>	69.4	67.0

(b) Crop size of local view.

Momentum	M	Seg.
0	69.2	66.9
0.1	69.0	67.1
0.9	70.5	68.1
0.99	68.7	66.6

(c) Momentum in EMA.

Table 5. Impact of hyper-parameters. The performance is evaluated on the VOC val set. The default settings are marked in gray.




Figure 7. Visualization of the attention map of class token w.r.t. patch tokens. The brighter region indicates a larger attention value. Left: the global view image in CTC; Right: the local view image randomly cropped from the global view in CTC.

CTC (ViT+PTC+CTC) and without CTC (ViT+PTC). Figure 5 also shows the proposed CTC can help to activate the less discriminative regions, which accounts for the quantitative improvements in Table 4.

A possible concern is that the objective of CTC may be limited when the global image includes multiple semantics but the local images only include partial semantics. However, as shown in Figure 7, the class token can capture multi-class semantic information. Therefore, when a local image only covers one semantic class, CTC can still enforce local-global representation consistency, though it's not a perfect objective. Moreover, considering the whole dataset, a semantic class usually co-occurs with multiple other classes, so the optimization to different classes can also be neutralized. Figure 5 shows CTC works well on images with multiple semantics.

**Fully-Supervised Counterparts.** The single-stage methods in Table 2 use different backbones. To ensure the fairness of comparison, we report their upper bound performance on the VOC val set, *i.e.* the performance of their fully-supervised counterparts. Table 3 shows that using only image-level labels as supervision, ToCo with ViT-B as backbone achieves 69.8% mIoU, which is 86.71% of its upper bound performance. In further, ToCo with ViT-B<sup>†</sup> as backbone can achieve higher performance. Compared to ToCo, the previous methods, AFA [34] and SLRNet [28] only achieve 83.86% and 83.17% of their fully-supervised counterparts, respectively. Particularly, WideResNet38 (WR38) can achieve comparable performance with ViT-B under full pixel-level supervision. However, in the WSSS experiments, ToCo can remarkably outperform SLRNet.

**Auxiliary Classifier.** Table 5a shows the impact of using different blocks to produce the auxiliary CAM. Usually, the shallow blocks cannot capture high-level semantics while the late blocks encounter the over-smoothing issue. For the ViT-B backbone used in our experiments (12 Transformer blocks), we empirically observe that appending the auxiliary classifier in the 10<sup>th</sup> block is a preferred choice.

**Local Crop Size.** The CTC module compares the representations of local and global view images to encourage consistency between the discriminative and less-discriminative object regions. Intuitively, a smaller local image may fail to include the foreground objects, mismatching the objective. On the contrary, a larger local image could contain too many discriminative object regions, also affecting distinguishing the uncertain regions. In Table 5b, we show that cropping local images with a size of 96<sup>2</sup> can achieve the best performance.

**EMA in CTC.** We use EMA to update the parameters of the global projection head  $\mathcal{P}^g$  in Figure 3. In Table 5c, we report the impact of the momentum value  $\rho$  in the EMA process. Table 5c shows  $\rho = 0.9$  is the best choice but other values can also yield favorable performance.

## 6. Conclusion

In this work, we aim to address the over-smoothing issue of ViT and further exploit its virtue for WSSS. Specifically, we first design a Patch Token Contrast module (PTC). PTC contrasts the final patch representations with knowledge extracted from the intermediate layer, which is empirically proved to counter the over-smoothing issue well. Inspired by the observation that the class token in ViT can capture the high-level semantics, we further propose a Class Token Contrast module (CTC) to contrast the class tokens of the local and global images, which can facilitate the representation consistency of the integral object regions. We plug ToCo into the single-stage WSSS framework and conduct extensive experiments on the VOC and COCO datasets. The experimental results show that ToCo can significantly outperform other competitors.

**Acknowledgements.** This work was supported in part by the National Natural Science Foundation of China under Grants 62225113 and 62002090, and the Science and Technology Major Project of Hubei Province (Next-Generation AI Technologies) under Grant 2019AEA170.



## References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*, pages 2209–2218, 2019. 1
- [2] 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
- [3] Nikita Araslanov and Stefan Roth. Single-stage semantic segmentation from image labels. In *CVPR*, pages 4253–4262, 2020. 5, 6, 7
- [4] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, pages 549–565. Springer, 2016. 1
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 2, 4, 5
- [7] 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 *CVPR*, pages 8991–9000, 2020. 2
- [8] 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. *TPAMI*, 40(4):834–848, 2017. 5
- [9] Qi Chen, Lingxiao Yang, Jianhuang Lai, and Xiaohua Xie. Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation. *CVPR*, 2022. 2
- [10] Tianlong Chen, Zhenyu Zhang, Yu Cheng, Ahmed Awadallah, and Zhangyang Wang. The principle of diversity: Training stronger vision transformers calls for reducing all levels of redundancy. In *CVPR*, pages 12020–12030, 2022. 4
- [11] Zhaozheng Chen, Tan Wang, Xiongwei Wu, Xian-Sheng Hua, Hanwang Zhang, and Qianru Sun. Class re-activation maps for weakly-supervised semantic segmentation. *CVPR*, 2022. 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 *ICLR*, 2020. 1, 2, 3, 5
- [13] Ye Du, Zehua Fu, Qingjie Liu, and Yunhong Wang. Weakly supervised semantic segmentation by pixel-to-prototype contrast. In *CVPR*, pages 4320–4329, 2022. 2, 6
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. 2, 5
- [15] Wei Gao, Fang Wan, Xingjia Pan, Zhiliang Peng, Qi Tian, Zhenjun Han, Bolei Zhou, and Qixiang Ye. Ts-cam: Token semantic coupled attention map for weakly supervised object localization. In *ICCV*, pages 2886–2895, October 2021. 1, 2, 4
- [16] Chengyue Gong, Dilin Wang, Meng Li, Vikas Chandra, and Qiang Liu. Vision transformers with patch diversification. *arXiv preprint arXiv:2104.12753*, 2021. 2, 4
- [17] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, pages 991–998, 2011. 5
- [18] Peng-Tao Jiang, Qibin Hou, Yang Cao, Ming-Ming Cheng, Yunchao Wei, and Hong-Kai Xiong. Integral object mining via online attention accumulation. In *ICCV*, pages 2070–2079, 2019. 2
- [19] 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 *CVPR*, pages 16886–16896, 2022. 1, 6
- [20] Meng Lan, Jing Zhang, Lefei Zhang, and Dacheng Tao. Learning to learn better for video object segmentation. *arXiv preprint arXiv:2212.02112*, 2022. 2
- [21] Jungbeom Lee, Jooyoung Choi, Jisoo Mok, and Sungroh Yoon. Reducing information bottleneck for weakly supervised semantic segmentation. *NeurIPS*, 34, 2021. 6
- [22] Jungbeom Lee, Seong Joon Oh, Sangdoon Yun, Junsuk Choe, Eunji Kim, and Sungroh Yoon. Weakly supervised semantic segmentation using out-of-distribution data. *CVPR*, 2022. 2, 6
- [23] Jungbeom Lee, Jihun Yi, Chaehun Shin, and Sungroh Yoon. Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In *CVPR*, pages 2643–2652, 2021. 1
- [24] 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. 2, 6
- [25] Jinlong Li, Zequn Jie, Xu Wang, Xiaolin Wei, and Lin Ma. Expansion and shrinkage of localization for weakly-supervised semantic segmentation. *NeurIPS*, 2022. 6
- [26] 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. Springer, 2014. 2, 5
- [27] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 5
- [28] Junwen Pan, Pengfei Zhu, Kaihua Zhang, Bing Cao, Yu Wang, Dingwen Zhang, Junwei Han, and Qinghua Hu. Learning self-supervised low-rank network for single-stage weakly and semi-supervised semantic segmentation. *IJCV*, 130(5):1181–1195, 2022. 5, 6, 7, 8
- [29] Namuk Park and Songkuk Kim. How do vision transformers work? In *ICLR*, 2021. 1, 3
- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 2

- [31] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lih Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021. 5, 6, 7
- [32] 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*, 2022. 5, 6
- [33] Lixiang Ru, Bo Du, Yibing Zhan, and Chen Wu. Weakly-supervised semantic segmentation with visual words learning and hybrid pooling. *IJCV*, 130(4):1127–1144, 2022. 2, 5, 6, 7
- [34] Lixiang Ru, Yibing Zhan, Baosheng Yu, and Bo Du. Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers. In *CVPR*, 2022. 1, 2, 3, 5, 6, 7, 8
- [35] 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
- [36] Han Shi, Jiahui Gao, Hang Xu, Xiaodan Liang, Zhenguo Li, Lingpeng Kong, Stephen MS Lee, and James Kwok. Revisiting over-smoothing in bert from the perspective of graph. In *ICLR*, 2021. 2
- [37] Yukun Su, Ruizhou Sun, Guosheng Lin, and Qingyao Wu. Context decoupling augmentation for weakly supervised semantic segmentation. In *ICCV*, pages 7004–7014, 2021. 2
- [38] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, pages 347–365. Springer, 2020. 2
- [39] Weixuan Sun, Jing Zhang, Zheyuan Liu, Yiran Zhong, and Nick Barnes. Getam: Gradient-weighted element-wise transformer attention map for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2112.02841*, 2021. 2
- [40] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *ECCV*, pages 507–522, 2018. 5
- [41] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357. PMLR, 2021. 6
- [42] Peihao Wang, Wenqing Zheng, Tianlong Chen, and Zhangyang Wang. Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. In *ICLR*, 2021. 1, 3
- [43] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, pages 12275–12284, 2020. 1, 2
- [44] 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. 2
- [45] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*, 2021. 2
- [46] Jinheng Xie, Xianxu Hou, Kai Ye, and Linlin Shen. Cross language image matching for weakly supervised semantic segmentation. *CVPR*, 2022. 2
- [47] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Multi-class token transformer for weakly supervised semantic segmentation. *CVPR*, 2022. 1, 2, 6
- [48] Yangyang Xu, Yibo Yang, and Lefei Zhang. Demt: Deformable mixer transformer for multi-task learning of dense prediction. *arXiv e-prints*, pages arXiv–2301, 2023. 2
- [49] 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 *CVPR*, pages 2623–2632, 2021. 2
- [50] 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. 5, 6
- [51] Bingfeng Zhang, Jimin Xiao, and Yao Zhao. Dynamic feature regularized loss for weakly supervised semantic segmentation. *arXiv preprint arXiv:2108.01296*, 2021. 1
- [52] Xiangrong Zhang, Zelin Peng, Peng Zhu, Tianyang Zhang, Chen Li, Huiyu Zhou, and Licheng Jiao. Adaptive affinity loss and erroneous pseudo-label refinement for weakly supervised semantic segmentation. In *ACM MM*, 2021. 5
- [53] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 1, 3
- [54] Tianfei Zhou, Meijie Zhang, Fang Zhao, and Jianwu Li. Regional semantic contrast and aggregation for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4299–4309, 2022. 2, 6