

# Weakly Supervised Semantic Segmentation by Pixel-to-Prototype Contrast

Ye Du, Zehua Fu, Qingjie Liu\*, Yunhong Wang

State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China  
Hangzhou Innovation Institute, Beihang University

{duyee, zehua\_fu, qingjie.liu, yhwang}@buaa.edu.cn

## Abstract

Though image-level weakly supervised semantic segmentation (WSSS) has achieved great progress with Class Activation Maps (CAMs) as the cornerstone, the large supervision gap between classification and segmentation still hampers the model to generate more complete and precise pseudo masks for segmentation. In this study, we propose weakly-supervised pixel-to-prototype contrast that can provide pixel-level supervisory signals to narrow the gap. Guided by two intuitive priors, our method is executed across different views and within per single view of an image, aiming to impose cross-view feature semantic consistency regularization and facilitate intra(inter)-class compactness(dispersion) of the feature space. Our method can be seamlessly incorporated into existing WSSS models without any changes to the base networks and does not incur any extra inference burden. Extensive experiments manifest that our method consistently improves two strong baselines by large margins, demonstrating the effectiveness. Specifically, built on top of SEAM, we improve the initial seed mIoU on PASCAL VOC 2012 from 55.4% to 61.5%. Moreover, armed with our method, we increase the segmentation mIoU of EPS from 70.8% to 73.6%, achieving new state-of-the-art.

## 1. Introduction

Benefiting from large-scale pixel-level annotations, semantic segmentation [38] has achieved remarkable progress in recent years. However, obtaining such precise pixel-wise annotations is laborious and time-consuming. To relieve this burden, many works resort to weakly supervised semantic segmentation (WSSS) that aims at learning segmentation models from weak labels such as image tags [2, 21, 22, 27, 32, 53, 62], bounding boxes [39], points [4] and scribbles [47]. Among them, image-level WSSS that requires only image tags has been extensively studied in the

\*Corresponding author.

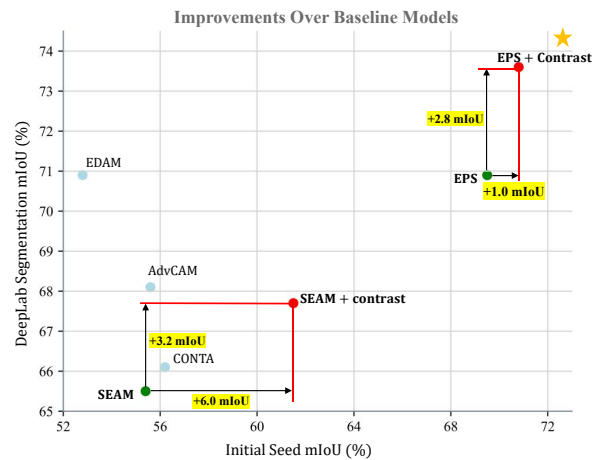


Figure 1. Initial seed quality vs. segmentation performance. Our method enables consistent performance improvements over state-of-the-arts, *i.e.*, SEAM [53] and EPS [32], without bringing any changes to the base networks during inference.

computer vision community.

Image-level WSSS is a challenging task since the image tags indicate only the existence of object categories and do not inform accurate object locations that are essential for semantic segmentation. To tackle this issue, Class Activation Maps (CAMs) [67] that identify which parts of the image contribute the most to the classification have been widely adopted to roughly estimate the regions of the target objects. The regions, also known as seeds, are used to generate pseudo ground truths for training segmentation models. However, CAMs only cover parts of objects, resulting in inaccurate and incomplete supervision. This issue is rooted in the supervision gap between the classification and the segmentation tasks. Specifically, classification networks supervised by image tags tend to focus on the most discriminative regions of objects for achieving a better performance of correctly tagging the image, while the segmentation task requires pixel-level supervision to assign a category to each pixel within the whole image. Narrowing down the super-

vision gap is crucial for WSSS, motivating us to explore pixel-wise supervisory signals that are complementary to image tags.

Inspired by the compelling contrastive self-supervised algorithms [20], we develop a novel *weakly-supervised pixel-to-prototype contrastive learning* method for WSSS, which can provide pixel-level supervision to improve the quality of CAMs and the performance of segmentation. Our method is based on two implicit but valuable priors: (i) features should retain semantic consistency across different views of an image; and (ii) pixels sharing the same label should have similar representations in the feature space, and vice versa. With these priors as guidelines, the pixel-to-prototype contrast is executed across different views and within per single view of each image, respectively, leading us to the cross-view contrast and intra-view contrast.

Our method is instantiated with a unified pixel-to-prototype contrastive learning formulation, which shapes the pixel embedding space through a prototype-based metric learning methodology. The core idea is pulling pixels together to their positive prototypes and pushing them away from their negative prototypes to learn discriminative dense visual representations. In our method, a prototype is defined as a representative embedding of a category. It is estimated from pixel-wise feature embeddings with the top activations in the CAMs. During learning, the polarity of each prototype is determined by the pseudo label of each pixel coupled to it in the current mini-batch. However, when generating pseudo masks from CAMs, there emerges a tricky problem: the over-activated and under-activated regions may corrupt the contrastive learning, especially the intra-view contrast. To alleviate this issue, we adopt two strategies: semi-hard prototype mining and hard pixel sampling to reduce the inaccurate contrasts as well as better utilize hard samples.

Recently, Wang *et al.* [53] propose SEAM to mitigate the supervision gap issue with a CAM equivariance constraint that enforces the CAMs to have the same spatial transformation as the input images. Our method has two major differences. Firstly, our method imposes regularizations to pixel-level features, enforcing the pixel embedding to be similar to the positive prototype and dissimilar to negative prototypes, while SEAM calculates consistency loss between CAMs of different views of the same image. Moreover, both cross-view and intra-view regularizations are considered in our work, while SEAM only integrates equivariant regularization across views.

Our method can be seamlessly incorporated into existing WSSS models without any changes to the base networks. It requires only additional common projectors during training and does not incur an extra inference burden. Experiments demonstrate that our method improves state-of-the-art models by large margins. As shown in Fig. 1, our method consistently improves two strong baseline models

*w.r.t.* both initial seed qualities and segmentation performance. We also validate our method through an extensive ablation study, where we find that each component contributes substantial performance improvements.

To summary, our main contributions are as follows:

- We propose weakly-supervised pixel-to-prototype contrast for WSSS. It enables the pixels to receive supervision from the reliable prototype of each class under the WSSS setting, which substantially narrows the gap between classification and segmentation.
- We propose to perform pixel-to-prototype contrastive learning both within a view and across different views of an image, which significantly improves the qualities of CAMs and the subsequent segmentation masks.
- Our method shows impressive results, surpassing baseline models by large margins and achieving the top performance on the standard benchmark.

## 2. Related Work

**Image-level WSSS.** WSSS with image-level labels has achieved significant progress under the pipeline of first generating pseudo masks then training a semantic segmentation network. Recent methods for WSSS rely on CAMs [67] to inform object locations by discovering image pixels that are informative for classification. However, CAMs can only highlight the most discriminative regions of objects, thus offering incomplete pseudo masks. Substantial efforts have been devoted to resolving this issue. They intend to complete the CAMs by enforcing networks to pay more attention to non-discriminative object regions using strategies such as region erasing [18, 54], region supervision [25], and region growing [19, 43]. Some other works refine the CAM by an iterative solution. For instance, PSA [2] and IRN [1] propose to propagate local responses to nearby areas which belong to the same semantic entity via random walk.

The inherent reason for the above issue is due to the supervision gap between classification and segmentation. Noticing this, many researchers explore to use additional supervision, such as multi-level feature maps [26], accumulated feature maps [21], cross-image semantics [15, 45], sub-categories [5], saliency maps [32, 60], and CAM consistency constraints [53], to narrow the gap. These approaches are simple but achieve encouraging performance.

**Contrastive Learning.** Contrastive learning (CL) [20] has shown great potential in learning discriminative representations without labels. The core idea of CL is to use the InfoNCE loss [40] to measure how well a model can classify a feature representation from a set of unrelated negative samples. For example, [58] learns feature representations at the instance level with a memory bank, where they try to maximally scatter instance embeddings over a unit sphere.

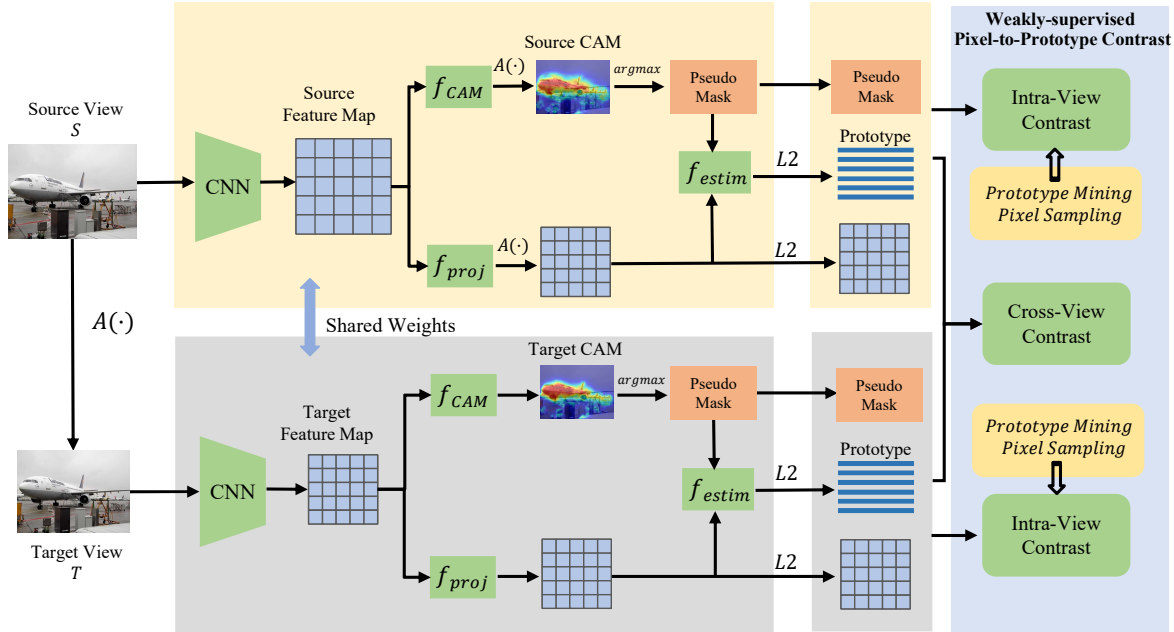


Figure 2. The overall pipeline of our proposed pixel-to-prototype contrast for WSSS.  $A(\cdot)$  is a spatial transformation for augmenting training samples.  $f_{CAM}$ ,  $f_{proj}$  are implemented by  $1 \times 1$  convolutional layer followed by ReLU.  $f_{est}$  represents the prototype estimation process and  $\mathcal{P}^{\{S,T\}}$  represent the generated prototypes.  $L2$  denotes per-pixel L2 normalization. The  $argmax$  function is conducted per-pixel along the channel dimension and returns the index of the maximum value.

MoCo [17] matches an encoded feature to a dynamic dictionary which evolves with a momentum updating strategy. SimCLR [9] presents a simple framework that engenders negative samples from large mini-batches.

Further, Khosla *et al.* [24] extend the self-supervised contrastive approach to the fully-supervised setting. A supervised contrastive loss is put forward to effectively leverage label information, enabling intra-class compactness and inter-class dispersion of the feature space. Wang *et al.* [52] propose dense contrastive learning that works at pixel-level and achieves superior performance than MoCo on downstream dense prediction tasks. In addition, Li *et al.* [33] propose to facilitate contrastive learning with clustering. They propose ProtoNCE loss that absorbs the advantages of contrastive learning and cluster-based unsupervised representation method [48], showing substantial improvements on several benchmarks.

**Contrastive Learning in Segmentation.** Recently, many studies leverage contrastive learning to promote image segmentation. These works utilize pixel-level or patch-level contrastive learning to improve semantic segmentation under fully supervised [50], semi-supervised [3], weakly-supervised [22], and unsupervised [49] settings. Caron *et al.* [22] improve WSSS by pixel-to-segment contrast, where they assume the segments are known in advance. In their work, they use SEAM to generate CAMs. Instead, we build

on top of existing WSSS methods and directly produce better quality seeds for segmentation. Recently, some studies apply contrastive learning to domain adaptation [35] and few-shot [36] semantic segmentation, also showing impressive results.

**Consistency Regularization.** Consistency regularization is a hot topic in the field of semi-supervised semantic segmentation. The idea is to enforce semantic or distribution consistency between various perturbations, such as image augmentation [23] and network perturbation [10, 64]. For example, Ke *et al.* [23] enforce cross probability consistency; Chen *et al.* [10] impose consistency regularization on two networks perturbed with different parameters for the same input image. The common goal of such methods is to construct appropriate supervisions by imposing consistency regularization, which greatly improves segmentation performance of semi-supervised segmentation. However, this idea is rarely studied in the case of weakly supervised segmentation. Our proposed cross-view pixel-to-prototype contrast can be regarded as imposing feature semantic consistency regularization across different views of each image.

### 3. Methodology

Our method can be interpreted as a regularization term that is adaptable to any existing WSSS framework without

changing the inference procedure. The overall loss function for training such a model is the linear combination of the cross-view contrastive loss  $\mathcal{L}^{\text{cross}}$  and the intra-view contrastive loss  $\mathcal{L}^{\text{intra}}$ :

$$\mathcal{L}^{\text{contrast}} = \alpha \mathcal{L}^{\text{cross}} + \beta \mathcal{L}^{\text{intra}} \quad (1)$$

where  $\alpha, \beta$  are two positive constants.

In this section, we first review how to generate CAMs, then introduce our proposed pixel-to-prototype contrast and how to estimate the prototypes, and finally elaborate on applying contrastive learning across different views and within per single view of each image. The framework of our method is illustrated in Fig. 2. Following the common practice, we first generate pixel-wise pseudo masks with our proposed method, then use them to train a DeepLab [6, 8] segmentation network.

### 3.1. Preliminary

We begin with a brief review on how CAMs are generated by visualization technique. Given a CNN (*e.g.* ResNet38 [57]), we denote the last convolutional feature maps by  $\mathbf{f} \in \mathbb{R}^{D \times HW}$ , where  $HW$  is the spatial size and  $D$  is the channel dimension. A Global Average Pooling (GAP) operation is applied to aggregate the feature maps. Next, a fully connected layer with parameters  $\mathbf{w} \in \mathbb{R}^{C \times D}$  is applied to retrieve the class scores. Here,  $C$  is the number of classes. Formally, the score for class  $c$  is obtained as

$$s_c = \frac{1}{HW} \sum_{j=1}^D \mathbf{w}_{c,j} \sum_i \mathbf{f}_{j,i} \quad (2)$$

After that, the CAM  $\mathbf{m}_c$  for class  $c$  is given by:

$$\mathbf{m}_c = \text{ReLU} \left( \sum_{j=1}^D \mathbf{w}_{c,j} \mathbf{f}_{j,:} \right) \quad (3)$$

It is worth noticing that a theoretically equivalent and much convenient way to compute CAMs is directly selecting the feature maps of the last convolutional layer [66]. Given the feature maps  $\mathbf{f}$ , we can add an additional convolutional layer which consists of  $C$  kernels with a size of  $1 \times 1$  and stride 1 on top of  $\mathbf{f}$  to obtain  $\mathbf{f}' \in \mathbb{R}^{C \times HW}$ . The score for class  $c$  is then computed by applying GAP to  $\mathbf{f}'$ . The maps  $\mathbf{f}'$  followed by ReLU function is directly used as the CAMs. In this paper, we follow this way for computing the CAMs.

### 3.2. Pixel-to-Prototype Contrast

Given the CAM of an image, we use a pixel-wise argmax function to generate the pseudo mask  $\mathbf{y}$ , *i.e.*,  $\mathbf{y} = \text{argmax}(\mathbf{m})$ , which determines the category of each pixel. For each category, there exists a representative embedding,




Figure 3. Illustration of the pixel-to-prototype contrast in the projected feature space. Pixel embeddings (circles) and prototypes (squares) of the same color are of the same category.

*i.e.*, the prototypes, denoted by  $\mathcal{P} = \{\mathbf{p}_c\}_{c=1}^C$ . Our goal is to learn discriminative feature embedding for each pixel aided by contrastive learning in a projected feature space. The idea is shown in Fig. 3. We first obtain pixel-wise projected feature  $\mathbf{v}_i \in \mathbb{R}^{128}$  by a projector, which is implemented with a  $1 \times 1$  convolutional layer followed by ReLU. Then, given  $\mathbf{v}_i$  and  $\mathcal{P}$ , the pixel-to-prototype contrast  $\mathcal{F}(\cdot)$  holds the following formulation:

$$\mathcal{F}(\mathbf{v}_i; \mathbf{y}_i; \mathcal{P}) = -\log \frac{\exp(\mathbf{v}_i \cdot \mathbf{p}_{\mathbf{y}_i} / \tau)}{\sum_{\mathbf{p}_c \in \mathcal{P}} \exp(\mathbf{v}_i \cdot \mathbf{p}_c / \tau)} \quad (4)$$

where  $\mathbf{y}_i \in [1, 2, \dots, C]$  is the pseudo label of pixel  $i$ , which determines the positive prototype  $\mathbf{p}_{\mathbf{y}_i}$ .  $\tau$  is the temperature parameter, which is set to 0.1 following a common practice.

### 3.3. Prototype Estimation

We further describe how we generate the prototypes. A possible solution is to mine the pixel-level semantic structure through clustering, as done in unsupervised semantic segmentation [11]. However, in the case of weakly supervised setting, this approach cannot make full use of image tag information and usually requires over-clustering to achieve better performance [5, 48]. The resulting clusters generally cannot well match the true categories.

In this work, we treat the pixel-wise CAM values as the confidences and propose to estimate the prototypes from pixel-wise feature embeddings that are with the highest confidences. Specifically, for all pixels assigned to class  $c$ , we empirically choose the ones with top  $K$  confidences to estimate the prototype. The prototype  $\mathbf{p}_c$  is computed as a weighted average of the projected pixel-wise embeddings:

$$\mathbf{p}_c = \frac{\sum_{i \in \Omega_c} \mathbf{m}_{c,i} \mathbf{v}_i}{\sum_{i' \in \Omega_c} \mathbf{m}_{c,i'}} \quad (5)$$

where  $\Omega_c$  is the collection of top  $K$  pixels of class  $c$ , and each pixel  $i$  has CAM value  $\mathbf{m}_{c,i}$ . A subsequent L2 nor-

malization is applied to each prototype. Here,  $K$  is a hyperparameter and a smaller  $K$  means higher confidence for computing the prototype.

Moreover, in order to capture the global context of the entire dataset, we compute prototypes across the training batch, that is, choosing pixels with the highest CAM values in the whole training batch.

### 3.4. Cross-view Contrast

Given the formulation of pixel-to-prototype contrastive loss shown in Eq. (4), we describe in detail how we apply cross-view contrast under the guidance of the cross-view semantic consistency. Specifically, given an image as the *source* view  $S$ , we generate a *target* view  $T$  by a spatial transformation  $A(\cdot)$ , as shown in Fig. 2. Then, the two views are encoded using a pre-trained CNN backbone, which are further processed to obtain two CAMs. We apply the same transformation  $A(\cdot)$  to the feature map and the CAM of the source view, with the same purpose as [53].

**Cross Prototype Contrast.** Considering that there should be semantic consistency between the two views, the prototype from one view can serve as supervisory signal to another view, and vice versa. Precisely, given a pixel  $i$  with its pseudo label  $\mathbf{y}_i \in [1, 2, \dots, C]$  and the projected feature embedding  $\mathbf{v}_i$ , the prototypes  $\mathcal{P}' = \{\mathbf{p}'_c\}_{c=1}^C$  from another view are borrowed to impose regularization to the current view. With the definition of pixel-to-prototype contrast shown in Eq. (4), the cross prototype contrastive loss is calculated by

$$\mathcal{L}^{\text{cp}} = \frac{1}{|I|} \sum_{i \in I} \mathcal{F}(\mathbf{v}_i; \mathbf{y}_i; \mathcal{P}') \quad (6)$$

where  $I$  denotes the whole image and  $|\cdot|$  represents the cardinality.

**Cross CAM Contrast.** Furthermore, the CAM from one view can be utilized to impose consistency regularization to another view as well. The CAM determines the pseudo mask of a view. Thus, for a pixel  $i$  with prototypes  $\mathcal{P}$  from its own view, we utilize the pseudo label  $\mathbf{y}'_i$  from another view to determine the positive prototype and negative prototypes. Analogously, the cross CAM contrastive loss can be written as:

$$\mathcal{L}^{\text{cc}} = \frac{1}{|I|} \sum_{i \in I} \mathcal{F}(\mathbf{v}_i; \mathbf{y}'_i; \mathcal{P}) \quad (7)$$

It is noticeable that the cross-view contrast is symmetrical, since both *source* view and *target* view can serve as the current view for computing  $\mathcal{L}^{\text{cp}}$  and  $\mathcal{L}^{\text{cc}}$ . Eventually, the respective  $\mathcal{L}^{\text{cp}}$  and  $\mathcal{L}^{\text{cc}}$  of the two views are added together as the total cross-view contrastive loss  $\mathcal{L}^{\text{cross}}$ . For simplicity, we only present the formulation of one view as follows:

$$\mathcal{L}^{\text{cross}} = \mathcal{L}^{\text{cp}} + \mathcal{L}^{\text{cc}} \quad (8)$$

### 3.5. Intra-view Contrast

**Intra-view Contrast.** According to the second hypothesis of intra-class compactness and inter-class dispersion, we further propose intra-view contrast that is conducted within per single view of each image. As opposed to the cross-view contrast, for a pixel  $i$  with pseudo label  $\mathbf{y}_i$ , the intra-view contrast takes prototypes  $\mathcal{P}$  from the current view to perform pixel-to-prototype contrastive learning:

$$\mathcal{L}^{\text{intra}} = \frac{1}{|I|} \sum_{i \in I} \mathcal{F}(\mathbf{v}_i; \mathbf{y}_i; \mathcal{P}) \quad (9)$$

The intra-view contrast is conducted on both views, and we do not list the symmetrical form in Eq. (9) for simplicity. However, we experimentally find that trivially introducing  $\mathcal{L}^{\text{intra}}$  could cause performance degeneration. The reason is that there are no precise pixel-wise annotations in the case of weakly supervised setting, the pseudo label  $\mathbf{y}_i$  assigned to pixel  $i$  could be inaccurate, resulting in inaccurate contrasts. Motivated by the hard example mining strategies in contrastive learning [42, 50], we alleviate this issue by introducing semi-hard prototype mining. Moreover, we also adopt a hard pixel sampling strategy to focus more on pixel samples hard for segmentation.

**Semi-hard Prototype Mining.** For a pixel  $i$ , the assigned label  $\mathbf{y}_i$  determines the positive prototype  $\mathbf{p}_{\mathbf{y}_i}$  and negative prototypes  $\mathcal{P}_N = \mathcal{P} \setminus \mathbf{p}_{\mathbf{y}_i}$ . Inspired by [50], rather than directly using  $\mathcal{P}_N$ , we adopt semi-hard prototype mining: for each pixel, we first collect the top 60% hardest negative prototypes, from which we choose 50% as the negative samples to compute the intra-view contrastive loss.

Here, a remaining question is how to define ‘harder’ prototypes. Following [50], for pixel  $i$ , we view the prototypes except  $\mathbf{p}_{\mathbf{y}_i}$  with dot products to pixel feature embedding  $\mathbf{v}_i$  closer to 1 to be harder, *i.e.*, prototypes that are similar to the pixel.

**Hard Pixel Sampling.** We also introduce hard pixel sampling to make better use of hard pixels. Specially, instead of using all pixels belonging to a prototype  $\mathbf{p}_c$  to calculate the intra-view contrastive loss, we adopt a per-class pixel sampling strategy: for each class, half of the pixels are randomly sampled and half are the hard ones.

Unlike [50], in this part, we define ‘harder’ pixels without ground truth during training. For a prototype  $\mathbf{p}_c$ , we view the belonging pixels with dot products to  $\mathbf{p}_c$  closer to  $-1$  to be harder, *i.e.*, pixels that are dissimilar to the prototype. The definition of ‘harder’ pixels is exactly the opposite of ‘harder’ prototypes, as the pixel far away from the corresponding prototype requires more attention to be pulled closer to the prototype in order to improve intra-class compactness.

We experimentally prove that equipped with the two strategies, we mitigate the effect of incorrect contrasts and

make better use of hard examples, which further improves performance.

## 4. Experiment

### 4.1. Datasets and Baselines

**Datasets.** We evaluate our proposed method on PASCAL VOC 2012 segmentation dataset [12], the standard benchmark for WSSS. The dataset consists of 21 classes including a background, with 1,464, 1,449, and 1,456 images for train, validation, and test set, respectively. Following the common practice in semantic segmentation, we use the augmented train set that consists of 10,582 images [16] for training. We report the mean Intersection-over-Union (mIoU) for evaluation, and the mIoU on the VOC test set is obtained from the official evaluation server.

**Baselines.** We choose two strong models, SEAM [53] and EPS [32] as our baselines. SEAM proposes a CAM equivariant regularization to narrow the supervision gap. EPS utilizes saliency maps as an additional supervision. They achieve state-of-the-art performance for WSSS. We build on top of these models to evaluate the effectiveness of our proposed method.

### 4.2. Implementation Details

Following SEAM and EPS, ResNet38 is adopted as the backbone network with output stride equals 8. The images are randomly rescaled to the range of [448, 768] by the longest edge and then cropped by  $448 \times 448$  as the input size of the network following [53]. We use rescale transformation that resizes the source image to a size of  $128 \times 128$ , keeping the multiple of the output stride. This is slightly different from SEAM, but the rescaling degree is nearly the same. The CNN backbone and projector share weights between the two views. The projected features used for contrastive loss have a dimension of 128. When imposing our proposed contrastive regularization  $\mathcal{L}^{contrast}$  to SEAM and EPS, we set  $\alpha = 0.1$  and  $\beta = 0.1$  in order to keep balance with classification loss. We follow the training and inference procedure in SEAM and EPS, including the training epoch, learning rate, learning rate decay policy, weight decay rate, and optimizer.

After generating the pseudo masks, we train three semantic segmentation networks to make fair comparisons with the baseline models. Concretely, DeepLab-LargeFOV with ResNet38 is trained to make comparison with SEAM; DeepLab-LargeFOV and DeepLab-ASPP with ResNet101 are trained to make comparison with EPS, respectively. During inference, we adopt multi-scale and flip operations as done in previous works. Standard dense CRF is used as a post-processing procedure to refine the final segmentation masks.

Table 1. Evaluation (mIoU (%)) of the initial seed (Seed), the seed with CRF (+CRF), and the pseudo mask (Mask) refined by PSA [2] on PASCAL VOC 2012 train set.

| Method                            | Venue   | Seed                 | +CRF                 | Mask                 |
|-----------------------------------|---------|----------------------|----------------------|----------------------|
| PSA [2]                           | CVPR'18 | 48.0                 | -                    | 61.0                 |
| Chang et al. [5]                  | CVPR'20 | 50.9                 | 55.3                 | 63.4                 |
| CONTA [62]                        | NIPS'20 | 56.2                 | 65.4                 | 66.1                 |
| EDAM [56]                         | CVPR'21 | 52.8                 | 58.2                 | 68.1                 |
| AdvCAM [31]                       | CVPR'21 | 55.6                 | 62.1                 | 68.0                 |
| ECS-Net [46]                      | ICCV'21 | 56.6                 | 58.6                 | -                    |
| OC-CSE [28]                       | ICCV'21 | 56.0                 | 62.8                 | 66.9                 |
| CPN [63]                          | ICCV'21 | 57.4                 | -                    | -                    |
| CDA [44]                          | ICCV'21 | 58.4                 | -                    | 66.4                 |
| <b>Improvement over baseline:</b> |         |                      |                      |                      |
| SEAM [53]                         | CVPR'20 | 55.4                 | 56.8                 | 63.6                 |
| Ours w/ SEAM                      | -       | 61.5 <sup>+6.1</sup> | 64.0 <sup>+7.2</sup> | 69.2 <sup>+5.6</sup> |
| EPS [32]                          | CVPR'21 | 69.5                 | 71.4                 | 71.6                 |
| Ours w/ EPS                       | -       | 70.5 <sup>+1.0</sup> | 73.3 <sup>+1.9</sup> | 73.3 <sup>+1.7</sup> |

Table 2. Evaluation (mIoU (%)) of the initial seed (Seed), the seed with CRF (+CRF), and the pseudo mask (Mask) refined by IRN [1] on PASCAL VOC 2012 train set.

| Method       | Venue   | Seed | +CRF | Mask |
|--------------|---------|------|------|------|
| IRN [1]      | CVPR'19 | 48.8 | 54.3 | 66.3 |
| MBMNet [37]  | MM'20   | 50.2 | -    | 66.8 |
| CONTA [62]   | NIPS'20 | 48.8 | -    | 67.9 |
| AdvCAM [31]  | CVPR'21 | 55.6 | 62.1 | 69.9 |
| Ours w/ SEAM | -       | 61.5 | 64.0 | 70.1 |

### 4.3. Seed and Pseudo Mask Evaluation

To verify the effectiveness of our method, we report the quantitative qualities of both initial seeds and pseudo masks on VOC. Following SEAM, the seeds are obtained by directly applying a range of thresholds to separate the foregrounds and backgrounds in the CAMs. The results are shown in Tab. 1. As can be seen, we improve SEAM by 6.0% and 7.2% mIoU on initial seed and seed+CRF, respectively, showing excellent performance. Figure 4 indicates that the CAMs generated by our method not only cover the target objects completely but also show accurate boundaries. Our generated CAMs are more accurate to match the ground truth segmentation masks than the baseline. Moreover, compared with the recent methods CONTA [62], CPN [63], and CDA [44] that built on top of SEAM, our method outperforms them by large margins. A considerable improvement is also observed when applying our method to EPS. Precisely, our method with EPS achieves 73.3% mIoU on seed+CRF, achieving state-of-the-art performance.

A typical pipeline in WSSS is refining the initial seed by region growing with a random walk strategy. Most methods refine their initial seeds with PSA [2] or IRN [1]. Therefore, we also compare the qualities of the refined pseudo masks




Figure 4. The visualization of CAMs. (a) Images. (b) GT masks. (c) CAMs produced by SEAM. (d) CAMs produced by our method. Our method generates better CAMs than SEAM in terms of both accuracy and completeness.

obtained by our method and other recent techniques. The 5th column of Tab. 1 shows the mIoU performance on the pseudo mask refined by PSA. It can be observed that our method surpasses existing methods by large margins. Note that the seed quality of EPS with our method is high, we do not further refine it with PSA. Instead, the seed+CRF is directly supplied as the pseudo mask.

Table 2 compares the performance between our method (with SEAM) and other counterparts refined using IRN. Our method performs remarkably better than the best performing counterpart AdvCAM [31], exceeding it by a large margin of 5.9% mIoU on the initial seed. As shown in the table, the pseudo mask performance gain of AdvCAM comes mainly from the refining procedure with IRN, while ours relies more on the high quality initial seed generated by our method. Nevertheless, our method achieves state-of-the-art performance under the IRN refinement settings.

#### 4.4. Segmentation Performance

Generally, the generated pseudo masks are used to train a semantic segmentation network in a fully supervised manner. To make a fair comparison, we report DeepLab-LargeFOV and Deeplab-ASPP segmentation performance of our method and compare them with existing methods in Tab. 3 and Tab. 4, respectively. The SEAM trains a DeepLab-LargeFOV network with ResNet38 as the backbone, achieving 64.5% and 64.7% mIoU on PASCAL VOC val and test sets. With the same settings except equipped with our method, we increase the segmentation mIoU by 3.2% and 1.7% on val and test sets, showing substantial improvement. Moreover, our method with SEAM even surpasses many models with more powerful backbones. With DeepLab-LargeFOV and ResNet101, EPS equipped with our method achieves 72.3% and 73.5% mIoU on PASCAL

Table 3. Segmentation performance (mIoU (%)) on Pascal VOC val and test sets using DeepLab-LargeFOV. The best result is marked in bold and the improvements over baseline model is marked in red. *S* means method using saliency maps.

| Method  | Backbone  | <i>S</i> | val                  | test                 |
|---|-----------|----------|----------------------|----------------------|
| SEC [27] <small>ECCV'16</small>               | VGG16     | ✓        | 50.7                 | 51.1                 |
| MDC [55] <small>CVPR'18</small>               | VGG16     | ✓        | 60.4                 | 60.8                 |
| MCOF [51] <small>CVPR'18</small>              | ResNet101 | ✓        | 60.3                 | 61.2                 |
| SeeNet [18] <small>NIPS'18</small>            | ResNet101 | ✓        | 63.1                 | 62.8                 |
| Lee <i>et al.</i> [30] <small>ICCV'19</small> | ResNet101 | ✓        | 66.5                 | 67.4                 |
| OAA+ [21] <small>ICCV'19</small>              | ResNet101 | ✓        | 65.2                 | 66.4                 |
| CIAN [15] <small>AAAI'20</small>              | ResNet101 | ✓        | 64.3                 | 65.3                 |
| MCIS [45] <small>ECCV'20</small>              | ResNet101 | ✓        | 66.2                 | 66.9                 |
| ICD [13] <small>CVPR'20</small>               | ResNet101 |          | 67.8                 | 68.0                 |
| ECS-Net [46] <small>ICCV'21</small>           | ResNet38  |          | 66.6                 | 67.6                 |
| Xu <i>et al.</i> [59] <small>ICCV'21</small>  | ResNet38  | ✓        | 69.0                 | 68.6                 |
| SEAM [53] <small>CVPR'20</small>              | ResNet38  |          | 64.5                 | 65.7                 |
| Ours w/ SEAM                                  | ResNet38  |          | 67.7 <sup>+3.2</sup> | 67.4 <sup>+1.7</sup> |
| EPS [32] <small>CVPR'21</small>               | ResNet101 | ✓        | 71.0                 | 71.8                 |
| Ours w/ EPS                                   | ResNet101 | ✓        | 72.3 <sup>+1.3</sup> | 73.5 <sup>+1.7</sup> |

Table 4. Segmentation performance (mIoU (%)) on Pascal VOC val and test sets using DeepLab-ASPP. The best result is marked in bold and the improvements over baseline model is marked in red. *S* means method using saliency maps.

| Method  | Backbone  | <i>S</i> | val                  | test                 |
|---|-----------|----------|----------------------|----------------------|
| PSA [2] <small>CVPR'18</small>                  | ResNet38  |          | 61.7                 | 63.2                 |
| IRN [1] <small>CVPR'19</small>                  | ResNet50  |          | 63.5                 | 64.8                 |
| FlickrNet [29] <small>CVPR'19</small>           | ResNet101 | ✓        | 64.9                 | 65.3                 |
| Zhang <i>et al.</i> [65] <small>ECCV'20</small> | ResNet50  | ✓        | 66.6                 | 66.7                 |
| Fan <i>et al.</i> [14] <small>ECCV'20</small>   | ResNet101 | ✓        | 67.2                 | 66.7                 |
| Chen <i>et al.</i> [7] <small>ECCV'20</small>   | ResNet101 |          | 65.7                 | 66.6                 |
| Chang <i>et al.</i> [5] <small>CVPR'20</small>  | ResNet101 |          | 66.1                 | 65.9                 |
| CONTA [62] <small>NIPS'20</small>               | ResNet101 |          | 66.1                 | 66.7                 |
| SPML [22] <small>ICLR'21</small>                | ResNet101 |          | 69.5                 | 71.6                 |
| AdvCAM [31] <small>CVPR'21</small>              | ResNet101 |          | 68.1                 | 68.0                 |
| EDAM [56] <small>CVPR'21</small>                | ResNet101 | ✓        | 70.9                 | 70.6                 |
| Yao <i>et al.</i> [61] <small>CVPR'21</small>   | ResNet101 | ✓        | 68.3                 | 68.5                 |
| DRS [25] <small>AAAI'21</small>                 | ResNet101 |          | 71.2                 | 71.4                 |
| Li <i>et al.</i> [34] <small>AAAI'21</small>    | ResNet101 | ✓        | 68.2                 | 68.5                 |
| WSGCN [41] <small>ICME'21</small>               | ResNet101 |          | 68.7                 | 69.3                 |
| CDA [44] <small>ICCV'21</small>                 | ResNet38  |          | 66.1                 | 66.8                 |
| CPN [63] <small>ICCV'21</small>                 | ResNet38  |          | 67.8                 | 68.5                 |
| EPS [32] <small>CVPR'21</small>                 | ResNet101 | ✓        | 70.9                 | 70.8                 |
| Ours w/ EPS                                     | ResNet101 | ✓        | 72.6 <sup>+1.7</sup> | 73.6 <sup>+2.8</sup> |

VOC val and test sets. In addition, we train a DeepLab-ASPP network with our generated pseudo masks. As shown in Tab. 4, our method outperforms all existing methods, achieving new state-of-the-art performance on PASCAL VOC 2012 benchmark. We present some segmentation results in Fig. 5, from which we can find that our method works well for images of both simple and challenging scenes.




Figure 5. Qualitative segmentation results on PASCAL VOC 2012 val set. (a) Images. (b) GT masks. (c) Segmentation masks predicted by DeepLab-ASPP (Ours w/ EPS).

Table 5. Ablation performance (mIoU %) of our proposed method.

| Method             | train        | train+CRF    | val          | val+CRF      |
|--------------------|--------------|--------------|--------------|--------------|
| CAM                | 47.43        | 52.40        | -            | -            |
| SEAM (baseline)    | 55.41        | 56.83        | 52.54        | 53.70        |
| <b>Ablations:</b>  |              |              |              |              |
| + Cross Prototype  | 59.08        | 61.98        | 55.75        | 58.33        |
| + Cross CAM        | 60.35        | 63.10        | 57.33        | 59.68        |
| + Intra View       | 59.80        | -            | 56.97        | -            |
| + Proto Mining     | 60.91        | 63.40        | 58.20        | 60.62        |
| + Pixel Sampling   | 61.54        | 64.05        | 58.41        | 60.81        |
| <b>Improvement</b> | <b>+6.13</b> | <b>+7.22</b> | <b>+5.87</b> | <b>+7.11</b> |

#### 4.5. Ablation Study

To analyse how each component in our proposed method helps to improve WSSS, we present extensive ablation studies in this section. Here, all experiments are done with SEAM on PASCAL VOC 2012 dataset.

**Effectiveness of each component.** First, we demonstrate the effectiveness of each component. The results are shown in Tab. 5. As can be seen, with the two cross-view semantic consistency regularization term, *i.e.* cross prototype contrast and cross CAM contrast, we improve the mIoU of SEAM from 55.41% to 60.35% on the train set, and from 52.54% to 57.35% on the val set. Furthermore, after applying the intra-view pixel-to-prototype contrast, we observe a slight performance drop on both train and val sets. We adopt two sample mining strategies to mitigate this issue, by which we increase the mIoU by 1.74% and 1.44%. Finally, with all these components, we increase the mIoU of SEAM by 6.13% and 5.87% on the train and val sets, showing significant improvements.

**Choices of  $K$ .** We empirically choose  $K$  pixel embeddings per-class in a mini-batch with the highest CAM values to evaluate the prototype. A smaller  $K$  indicates to use more confident pixels for estimation, while a larger  $K$  could

Table 6. Effects of  $K$  for estimating the prototypes of our method. The results (mIoU (%)) are reported on PASCAL VOC 2012 train set.  $K=0$  means the SEAM baseline without contrast.

| $K$  | 0     | 4     | 8     | 16    | 32           | 64    |
|------|-------|-------|-------|-------|--------------|-------|
| mIoU | 55.41 | 60.02 | 60.87 | 61.08 | <b>61.54</b> | 60.55 |

Table 7. Ablations for more spatial transformations. The results (mIoU (%)) are reported on PASCAL VOC 2012 train set.

| Rescale | Flip | Rotation | Translation | mIoU         |
|---------|------|----------|-------------|--------------|
|         |      |          |             | 47.43        |
| ✓       |      |          |             | 61.54        |
| ✓       | ✓    |          |             | <b>61.63</b> |
| ✓       |      | ✓        |             | 58.57        |
| ✓       |      |          | ✓           | 59.31        |

probably incur some wrong pixels from other categories. We conduct experiments to analyse how  $K$  affects the initial seed performance. As shown in Tab. 6, we test a wide range of  $K$ . The results show that our proposed method is robust to the selection of  $K$ . Eventually, we set  $K = 32$  with the best performance.

**Spatial transformation.** Similar to SEAM, the transformation in our method can be any spatial transformation. In our implementation, we rescale the source image of size  $448 \times 448$  to  $128 \times 128$  to get the target view in order to make sure the input size of the target view is a multiple of the output stride. This setting is slightly different from SEAM but remains almost identical to it. Following SEAM, we also adopt three transformations to evaluate our method: random rotation in  $[-20, 20]$  degrees, translation by 16 pixels and horizontal flip. The results are shown in Tab. 7. It can be seen that simply incorporating different transformations is not much effective, a similar observation to SEAM.

## 5. Conclusion

In this paper, we propose weakly-supervised pixel-to-prototype contrast that provides pixel-level supervisory signals to narrow the supervision gap and improve image-level WSSS. The pixel-to-prototype contrast is performed on both cross-view and intra-view of an image, which imposes cross-view feature semantic consistency regularization and facilitates intra(inter)-class compactness(dispersion) of the feature space. Extensive experiments validate the superiority of our method. In future, we will explore more prototype estimation ways and more application scenarios of our method, such as fully-supervised and semi-supervised segmentation.

**Acknowledgement.** The work was supported by the National Key Research and Development Program of China under Grant 2018YFB1701600, National Natural Science Foundation of China under Grant U20B2069.



## References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*, 2019. 2, 6, 7
- [2] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, 2018. 1, 2, 6, 7
- [3] Iñigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C. Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *ICCV*, 2021. 3
- [4] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 1
- [5] 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*, 2020. 2, 4, 6, 7
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015. 4
- [7] Liyi Chen, Weiwei Wu, Chenchen Fu, Xiao Han, and Yuntao Zhang. Weakly supervised semantic segmentation with boundary exploration. In *ECCV*, 2020. 7
- [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*, 2017. 4
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3
- [10] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *CVPR*, 2021. 3
- [11] Jang Hyun Cho. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *CVPR*, 2021. 4
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 6
- [13] Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In *CVPR*, 2020. 7
- [14] Junsong Fan, Zhaoxiang Zhang, and Tieniu Tan. Employing multi-estimations for weakly-supervised semantic segmentation. In *ECCV*, 2020. 7
- [15] Junsong Fan, Zhaoxiang Zhang, Tieniu Tan, Chunfeng Song, and Jun Xiao. Cian: Cross-image affinity net for weakly supervised semantic segmentation. In *AAAI*, 2020. 2, 7
- [16] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 6
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3
- [18] Qibin Hou, Peng-Tao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *NIPS*, 2018. 2, 7
- [19] Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *CVPR*, 2018. 2
- [20] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 2021. 2
- [21] 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*, 2019. 1, 2, 7
- [22] Tsung-Wei Ke, Jyh-Jing Hwang, and Stella Yu. Universal weakly supervised segmentation by pixel-to-segment contrastive learning. In *ICLR*, 2021. 1, 3, 7
- [23] Zhanghan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *ECCV*, 2020. 3
- [24] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *NIPS*, 2020. 3
- [25] Beomyoung Kim, Sangeun Han, and Junmo Kim. Discriminative region suppression for weakly-supervised semantic segmentation. In *AAAI*, 2021. 2, 7
- [26] Hyo-Eun Kim and Sangheum Hwang. Deconvolutional feature stacking for weakly-supervised semantic segmentation. *arXiv:1602.04984*, 2016. 2
- [27] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *ECCV*, 2016. 1, 7
- [28] Hyeokjun Kweon, Sung-Hoon Yoon, Hyeonseong Kim, Daehee Park, and Kuk-Jin Yoon. Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation. In *ICCV*, 2021. 6
- [29] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *CVPR*, 2019. 7
- [30] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Frame-to-frame aggregation of active regions in web videos for weakly supervised semantic segmentation. In *ICCV*, 2019. 7
- [31] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *CVPR*, 2021. 6, 7
- [32] 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*, 2021. 1, 2, 6, 7
- [33] Junnan Li, Pan Zhou, Caiming Xiong, and Steven C. H. Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021. 3

- [34] Xueyi Li, Tianfei Zhou, Jianwu Li, Yi Zhou, and Zhaoxiang Zhang. Group-wise semantic mining for weakly supervised semantic segmentation. In *AAAI*, 2021. 7
- [35] Weizhe Liu, David Ferstl, Samuel Schuster, Lukas Zebedin, Pascal Fua, and Christian Leistner. Domain adaptation for semantic segmentation via patch-wise contrastive learning. *arXiv:2104.11056*, 2021. 3
- [36] Weide Liu, Zhonghua Wu, Henghui Ding, Fayao Liu, Jie Lin, and Guosheng Lin. Few-shot segmentation with global and local contrastive learning. *arXiv:2108.05293*, 2021. 3
- [37] Weide Liu, Chi Zhang, Guosheng Lin, Tzu-Yi Hung, and Chunyan Miao. Weakly supervised segmentation with maximum bipartite graph matching. In *MM*, 2020. 6
- [38] Shervin Minaee, Yuri Y Boykov, Fatih Porikli, Antonio J Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *TPAMI*, 2021. 1
- [39] Youngmin Oh, Beomjun Kim, and Bumsub Ham. Background-aware pooling and noise-aware loss for weakly-supervised semantic segmentation. In *CVPR*, 2021. 1
- [40] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. 2
- [41] Shun-Yi Pan, Cheng-You Lu, Shih-Po Lee, and Wen-Hsiao Peng. Weakly-supervised image semantic segmentation using graph convolutional networks. In *ICME*, 2021. 7
- [42] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. In *ICLR*, 2021. 5
- [43] Wataru Shimoda and Keiji Yanai. Self-supervised difference detection for weakly-supervised semantic segmentation. In *ICCV*, 2019. 2
- [44] Yukun Su, Ruizhou Sun, Guosheng Lin, and Qingyao Wu. Context decoupling augmentation for weakly supervised semantic segmentation. In *ICCV*, 2021. 6, 7
- [45] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, 2020. 2, 7
- [46] Kunyang Sun, Haoqing Shi, Zhengming Zhang, and Yongming Huang. Ecs-net: Improving weakly supervised semantic segmentation by using connections between class activation maps. In *ICCV*, 2021. 6, 7
- [47] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *ECCV*, 2018. 1
- [48] Kai Tian, Shuigeng Zhou, and Jihong Guan. Deepcluster: A general clustering framework based on deep learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2017. 3, 4
- [49] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. In *ICCV*, 2021. 3
- [50] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. In *ICCV*, 2021. 3, 5
- [51] Xiang Wang, Shaodi You, Xi Li, and Huimin Ma. Weakly-supervised semantic segmentation by iteratively mining common object features. In *CVPR*, 2018. 7
- [52] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, 2021. 3
- [53] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, 2020. 1, 2, 5, 6, 7
- [54] 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*, 2017. 2
- [55] Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *CVPR*, 2018. 7
- [56] 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 *CVPR*, 2021. 6, 7
- [57] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 2019. 4
- [58] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 2
- [59] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, Ferdous Sohel, and Dan Xu. Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation. In *ICCV*, 2021. 7
- [60] Qi Yao and Xiaojin Gong. Saliency guided self-attention network for weakly and semi-supervised semantic segmentation. *IEEE Access*, 2020. 2
- [61] 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*, 2021. 7
- [62] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xian-Sheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *NIPS*, 2020. 1, 6, 7
- [63] Fei Zhang, Chaochen Gu, Chenyue Zhang, and Yuchao Dai. Complementary patch for weakly supervised semantic segmentation. In *ICCV*, 2021. 6, 7
- [64] Liheng Zhang and Guo-Jun Qi. Wcp: Worst-case perturbations for semi-supervised deep learning. In *CVPR*, 2020. 3
- [65] Tianyi Zhang, Guosheng Lin, Weide Liu, Jianfei Cai, and Alex Kot. Splitting vs. merging: Mining object regions with discrepancy and intersection loss for weakly supervised semantic segmentation. In *ECCV*, 2020. 7
- [66] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *CVPR*, 2018. 4
- [67] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 1, 2