

Self Correspondence Distillation for End-to-End Weakly-Supervised Semantic Segmentation

Rongtao Xu^{1,3,*}, Changwei Wang^{1,3,*}, Jiayi Sun^{1,3},
Shibiao Xu^{2,†}, Weiliang Meng^{1,3,†}, Xiaopeng Zhang^{1,3}

¹NLPR, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, Beijing University of Posts and Telecommunications

³School of Artificial Intelligence, University of Chinese Academy of Sciences

shibiaoxu@bupt.edu.cn, {xurongtao2019, wangchangwei2019, sunjiayi2020, weiliang.meng, xiaopeng.zhang}@ia.ac.cn

Abstract

Efficiently training accurate deep models for weakly supervised semantic segmentation (WSSS) with image-level labels is challenging and important. Recently, end-to-end WSSS methods have become the focus of research due to their high training efficiency. However, current methods suffer from insufficient extraction of comprehensive semantic information, resulting in low-quality pseudo-labels and sub-optimal solutions for end-to-end WSSS. To this end, we propose a simple and novel Self Correspondence Distillation (SCD) method to refine pseudo-labels without introducing external supervision. Our SCD enables the network to utilize feature correspondence derived from itself as a distillation target, which can enhance the network’s feature learning process by complementing semantic information. In addition, to further improve the segmentation accuracy, we design a Variation-aware Refine Module to enhance the local consistency of pseudo-labels by computing pixel-level variation. Finally, we present an efficient end-to-end Transformer-based framework (TSCD) via SCD and Variation-aware Refine Module for the accurate WSSS task. Extensive experiments on the PASCAL VOC 2012 and MS COCO 2014 datasets demonstrate that our method significantly outperforms other state-of-the-art methods. Our code is available at <https://github.com/Rongtao-Xu/RepresentationLearning/tree/main/SCD-AAAI2023>.

Introduction

Weakly Supervised Semantic Segmentation (WSSS) aims to use weak/cheap labels to alleviate the reliance on time-consuming pixel-level annotations (Zhang et al. 2021). The most challenging scenario for WSSS methods is to use only image-level labels which we focus on, as it has no real localization information, and usually relies on multi-stage solutions to generate high-quality pseudo-labels.

Multi-stage methods (Hou et al. 2018; Lee et al. 2019; Wu et al. 2021) first train a classification model and then train a semantic segmentation network with refined pseudo-labels, which are generated from class activation maps (Zhou et al.




Figure 1: Visualization examples of CAMs generated by our TSCD. The corresponding categories are drawn on the original image with red crosses.

2016) (CAMs). Due to the complex pipeline of multi-stage methods and the need to train multiple models, recently some end-to-end methods (Araslanov and Roth 2020; Zhang et al. 2020a, 2021; Ru et al. 2022) have been proposed to speed up the training process. However, these methods rely heavily on the pseudo-labels generated by CAM as supervision information and still show limited performance in obtaining the comprehensiveness of semantic information and segmentation accuracy. The recent work (Caron et al. 2021) reveals that there is corresponding information about image semantic segmentation in the dense features of self-supervised transformers. We find that these dense features can be class activation maps, and further construct novel CAM feature correspondence to learn semantic correspondence information. In this paper, we propose a simple and novel method named Self Correspondence Distillation (SCD), which can enhance the network feature learning process by complementing semantic information to refine CAM without external supervision. Our SCD allows the network to utilize the CAM feature correspondence derived from itself as the distillation target for segmentation prediction features. This mechanism can help the network obtain comprehensive image semantic information and improve accuracy.

To further refine the initial pseudo-labels obtained by CAM, we design a Variation-aware Refine Module (VARM), which introduces the idea of image noise reduc-

*These authors contributed equally.

†Shibiao Xu and Weiliang Meng are the corresponding authors. Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tion, and updates pseudo-labels by calculating image pixel-level variation and image local information. Based on SCD and VARM, we provide an end-to-end Transformer-based framework with Self Correspondence Distillation (TSCD) for WSSS. We also visualize the CAM generated by our TSCD in Figure 1, and the results show that our TSCD can generate fine-grained CAMs by capturing comprehensive image semantic information.

In summary, our main contributions are as follows:

- We propose a novel Self Correspondence Distillation method, namely SCD, to enhance the feature learning of WSSS networks. According to our knowledge, this is the first attempt to use the network’s own CAM feature correspondence as the distillation target.
- We design the Variation-aware Refine Module (VARM) to address the local inconsistency of pseudo-labels. Our VARM refines the CAM by computing image pixel-level variation and employing pixel-adaptive convolution.
- We provide an end-to-end Transformer-based framework (TSCD) with Self Correspondence Distillation and VARM for weakly supervised semantic segmentation. Our TSCD achieves state-of-the-art segmentation performance on PASCAL VOC 2012 and MS COCO 2014.

Related Work

Weakly-Supervised Semantic Segmentation

Recently, deep learning has been developed rapidly (Wang et al. 2022b,d; Xu et al. 2022b, 2023; Zhao et al. 2021), especially in image segmentation (Xu et al. 2021b, 2022a; Wang et al. 2022a, 2021, 2022c; Dong et al. 2020). Weakly-Supervised Semantic Segmentation (WSSS) methods are mainly divided into two types: Multi-stage Methods and End-to-End Methods, both of which rely on class activation maps (Zhou et al. 2016) to generate pseudo labels. For Multi-stage Methods, the erasure methods (Hou et al. 2018) make the classifier focus on the complete object region by erasing the discriminative region. To obtain multiple attribute maps of target objects, Lee et al. (Lee et al. 2019) utilize dilated convolutions to accumulate activation regions, while (Wu et al. 2021) explore semantic similarities and differences across multiple input images. Several recent approaches attempt to explain and improve CAM generation, such as (Zhang et al. 2020b) from the perspective of causal relationships between images and labels, and (Lee et al. 2021a) from the perspective of information bottleneck theory. Due to the rough boundary of the obtained initial localization map, refinement techniques such as CRF and IRN (Ahn, Cho, and Kwak 2019) are applied for subsequent refinements.

While most WSSS methods are based on multi-stage, recent End-to-End Methods have also achieved comparable performance to multi-stage methods. For End-to-End Methods, training a well-performing end-to-end model is challenging due to the extremely limited supervision brought by the rough initial CAM. To address this problem, many studies have focused on improving pseudo-segmentation labels obtained from CAM maps. Araslanov et al. (Araslanov and

Roth 2020) proposed 1Stage which achieved great progress on the WSSS task via pixel-adaptive mask refinement, while Zhang et al. (Zhang et al. 2020a) used CRF on CAM to generate fine labels as supervision. Besides, Zhang et al. (Zhang et al. 2021) introduce feature-to-prototype alignment loss and adaptive affinity field, and Ru et al. (Ru et al. 2022) learn affinity from attention to refine CAM. Different from the above methods, we refine the CAM by learning the feature correspondence of the network’s own CAMs.

Knowledge Distillation and Self-Supervised Feature Learning

Knowledge distillation (Hinton, Vinyals, and Dean 2015) is a model compression method by transferring knowledge from large models to small models. Usually the student network and the teacher network share the same capacity, and the student network imitates the intermediate output of the teacher network. Recent studies (Hou et al. 2019) focus on attention distillation, where the student network learns attention maps obtained from the teacher network. Hou et al. (Hou et al. 2019) proposed SAD that can learn attention knowledge without a teacher network. Different from (Hou et al. 2019), we exploit feature correspondence to obtain comprehensive image semantic information without teacher network, inspired by self-supervised feature learning.

The goal of self-supervised feature learning is to learn meaningful visual features without human annotation, while contrastive learning over multiple augmentations is used to address this problem (Chen et al. 2020). Recently, Pinheiro et al. (O Pinheiro et al. 2020) generate spatially dense feature maps to compare local pixel features. (Hamilton et al. 2022) treats pretrained self-supervised features as input. Different from the above methods, our TSCD refines the CAM without external supervision, by taking the CAM feature correspondence as the distillation target, and our novel idea is that the CAM feature correspondence may be extracted for self-learning.

Proposed Method

Preliminary

The Class Activation Map (CAM) was first proposed in (Zhou et al. 2016) to enable classifiers learn object localization. Given a feature map $f \in \mathbb{R}^{H \times W \times D}$ extracted by CNN and the class c , we apply global average pooling and fully connected layer operations to compute class scores:

$$y_c = \frac{1}{HW} \sum_k w_{c,k} \sum_i f_{k,i} \quad (1)$$

where $w_{c,k}$ represents the parameters of the fully connected layer. CAM is achieved by weighting the contribution of each channel of the feature map to a specific category of the classification score. Formally, the class activation map m_c of the class c is:

$$m_c = Relu(\sum_k w_{c,k} f_{k,:}). \quad (2)$$

In this paper, we adopt the class activation map as the initial pseudo-label.

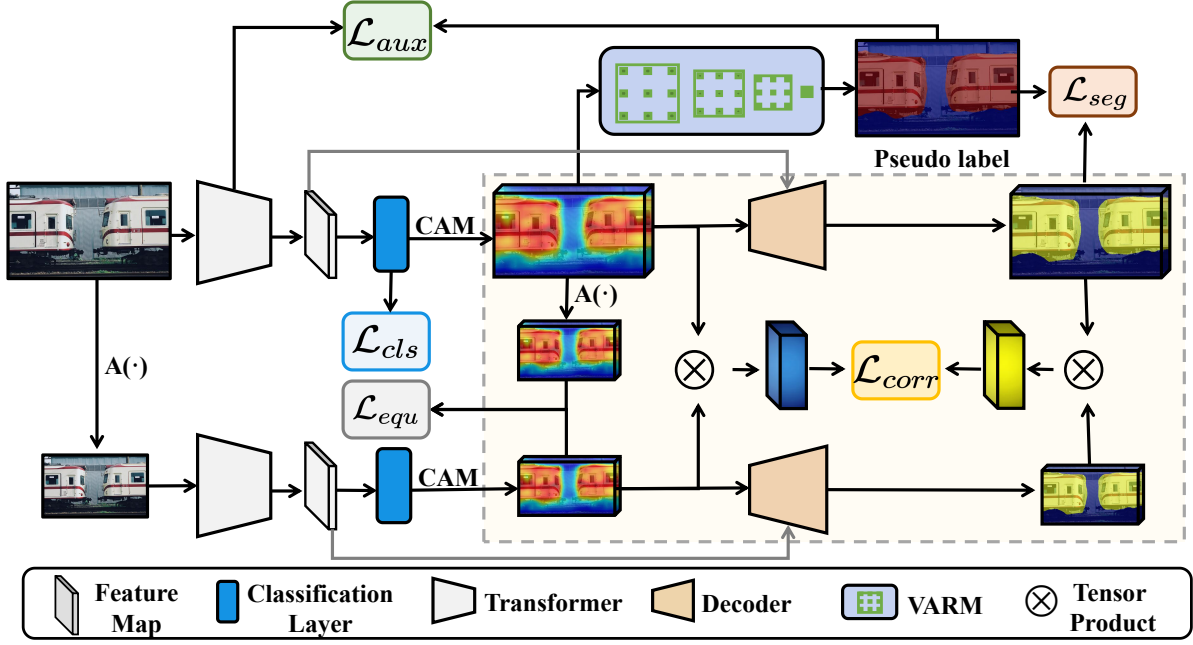


Figure 2: Illustration of the proposed end-to-end framework TSCD. Our TSCD uses the Transformer backbone as the encoder and adopts CAM (Zhou et al. 2016) to generate initial pseudo-labels. The generated initial pseudo-labels are then refined with our SCD and VARM. The yellow area indicates the application of our SCD. The optimization of the network consists of self correspondence distillation loss, equivariant regularization loss, classification loss, and auxiliary loss, etc. $A(\cdot)$ means affine transformation.

Self Correspondence Distillation

Our goal is to perform Self Correspondence Distillation to refine the CAM of the original image. Our Self Correspondence Distillation does not require any additional labels or external supervision, while it can assist the network to obtain comprehensive image semantic information. Recent advances in self-supervised feature learning have shown that dense features are semantically related (Collins, Achanta, and Susstrunk 2018; Zhou et al. 2016). In this paper, we verify that these dense feature maps can be class activation maps, and further construct CAM feature correspondence for feature learning.

CAM Feature Correspondence: Formally, we focus on the correlation volume between CAMs. Given two CAMs $m_1 \in \mathbb{R}^{H_1 \times W_1 \times C}$ and $m_2 \in \mathbb{R}^{H_2 \times W_2 \times C}$, with H_1, H_2 for height, W_1, W_2 for width, and C for category, we define the CAM feature correspondence as:

$$\mathcal{M}_{h_1 w_1 h_2 w_2} = \sum_c \frac{m_1 \cdot m_2}{|m_1| |m_2|}. \quad (3)$$

Specifically, given an image $I \in \mathbb{R}^{H \times W \times D}$ and an affine transformation $A(\cdot)$, we employ the feature maps extracted by the encoder to generate CAMs. Taking CAM m_1 as an example, we use $E : \mathbb{R}^{H \times W \times D} \rightarrow \mathbb{R}^{H_1 \times W_1 \times C}$ to denote the process of CAM m_1 generation from the extracted feature maps. Therefore, m_1 and m_2 can be expressed as:

$$\begin{cases} m_1 = E(I) \\ m_2 = E(A(I)) \end{cases} \quad (4)$$

Then the CAM feature correspondence can be further expressed as:

$$\mathcal{M}_{h_1 w_1 h_2 w_2} = \sum_c \frac{E(I) \cdot E(A(I))}{|E(I)| |E(A(I))|} \quad (5)$$

whose entries represent the cosine similarity between the feature at m_1 's position (h_1, w_1) and m_2 's position (h_2, w_2) . As shown in Figure 2, the encoder and decoder share weights, and the segmentation prediction maps are denoted as $s_1 \in \mathbb{R}^{H_1 \times W_1 \times C}$ and $s_2 \in \mathbb{R}^{H_2 \times W_2 \times C}$ for images I and $A(I)$, respectively. Then the segmentation feature correspondence is defined as:

$$\mathcal{S}_{h_1 w_1 h_2 w_2} = \sum_c \frac{s_1 \cdot s_2}{|s_1| |s_2|}. \quad (6)$$

SCD for Training: The intuition behind our Self Correspondence Distillation (SCD) is that segmentation feature correspondence can extract useful semantic information from the CAM feature correspondence to refine CAMs in a self-learning manner. Inspired by the self-supervised feature learning, we consider aligning the segmentation feature correspondence with the network's own CAM feature correspondence. The loss function is designed to push the corresponding CAMs together to enhance semantic relatedness,

when there is a significant correlation between two segmentation predictions. We implement the SCD loss function by simple element-wise multiplication of the segmentation feature corresponding $\mathcal{S}_{h_1 w_1 h_2 w_2}$ and the CAM feature corresponding $\mathcal{M}_{h_1 w_1 h_2 w_2}$:

$$\mathcal{L}_{SCD} = - \sum_{h_1 w_1 h_2 w_2} \mathcal{M}_{h_1 w_1 h_2 w_2} \cdot \max(\mathcal{S}_{h_1 w_1 h_2 w_2}, 0). \quad (7)$$

where $\max(\cdot, 0)$ means the zero clamping. In practice, in order to ensure the inference efficiency, we adopt a random sampling strategy to train our SCD loss function, and the number of samples is n . If the size of the segmentation prediction map is different from the size of the corresponding CAM, then bilinear upsampling is applied to the segmentation prediction map.

Variation-Aware Refine Module

The initial pseudo-labels obtained by CAM are usually locally inconsistent and have rough boundaries. Many multi-stage methods employ CRF to further refine the initial pseudo-labels, which reduces the training efficiency. For end-to-end methods, Araslanov et al. (Araslanov and Roth 2020) utilize pixel-adaptive convolution (Su et al. 2019) to extract local image information for local consistency, while Ru et al. (Ru et al. 2022) further combine spatial information to build a refinement module. Different from (Araslanov and Roth 2020; Ru et al. 2022), we design the Variation-aware Refine Module, which introduces the idea of image noise reduction to overcome local inconsistencies.

Specifically, for positions (i, j) and (k, l) in image I , we first compute the image pixel-level variation:

$$\mathcal{V}_{ij,kl} = (x_{i,j-1,kl} - x_{i,j,kl})^2 + (x_{i+1,j,kl} - x_{i,j,kl})^2. \quad (8)$$

Next we calculate the local information term $k_{ij,kl}^{rgb}$:

$$k_{ij,kl}^{rgb} = \frac{-(\alpha |I_{i,j} - I_{k,l}|)^2}{\sigma_{ij}^2} \quad (9)$$

where σ_{ij} represents the standard deviation and α represents the smoothing weight. To enhance the local consistency of pseudo-labels, for pixels with large variation in the image, we calculate the correction kernel $k_{ij,kl}$ to avoid some abruptly deformed values:

$$k_{ij,kl} = \frac{\exp(k_{ij,kl}^{rgb})}{\sum_{(x,y) \in N(i,j)} \exp(k_{ij,kl}^{rgb})} - \beta \frac{\exp(\mathcal{V}_{ij,kl})}{\sum_{(x,y) \in N(i,j)} \exp(\mathcal{V}_{ij,kl})} \quad (10)$$

where $N(i, j)$ is the set of neighbors of (i, j) , obtained using dilated convolution as defined in (Araslanov and Roth 2020). We adopt an iterative update strategy to update the pixel label (CAM) $\mathcal{P}^{i,j,c}$:

$$\mathcal{P}_t^{i,j,c} = \sum_{(k,l) \in N(i,j)} k_{ij,kl} \mathcal{P}_{t-1}^{k,l,c}. \quad (11)$$

Our Variation-aware Refine Module enhances the local consistency of the initial pseudo-labels by perceiving pixel-level variation, while ensuring high training efficiency.

Transformer-Based Framework with Self Correspondence Distillation

As shown in Figure 2, our Transformer-based framework with Self Correspondence Distillation (TSCD) consists of a transformer backbone, SCD, VARM, equivariant regularization loss, classification loss, auxiliary loss and segmentation loss. Next, each loss function and total loss are introduced separately.

Equivariant Regularization Loss: Equivariant constraints have been shown to narrow the supervision gap between weak supervision and full supervision (Du et al. 2021; Wang et al. 2020). We construct an equivariant constraint using an equivariant regularization loss:

$$\mathcal{L}_{equ} = \|A(m_1) - m_2\|_1 = \|A(E(I)) - E(A(I))\|_1. \quad (12)$$

Classification Loss: As shown in Figure 2, the aggregated feature maps are fed into the classification layer to compute prediction vector p for the image-level classification. We adopt the multi-label soft-margin loss as the classification loss for network training. For the total number of classes C , the classification loss is defined as:

$$\mathcal{L}_{cls} = \frac{1}{C} \sum_{c=1}^C (l^c \log(\frac{1}{1 + e^{-p_c}}) + (1 - l^c) \log(\frac{e^{-p_c}}{1 + e^{-p_c}})) \quad (13)$$

where l is the ground truth for class labels.

Auxiliary Loss: To further enhance the performance of our network, we adopt the affinity loss of (Ru et al. 2022) as our auxiliary loss in Figure 2. Affinity loss is beneficial to the self-attention learning of the encoder and helps the network focus on complete object regions. Different from (Ru et al. 2022), we directly use the attention maps $(\mathcal{A}_1, \mathcal{A}_2)$ output by the last two layers of the encoder to calculate the auxiliary loss. Formally, the auxiliary loss is expressed as:

$$\mathcal{L}_{aux} = \frac{1}{N^-} \sum_{R^-} (1 - \frac{1}{1 + e^{-\text{concat}(\mathcal{A}_1, \mathcal{A}_2)}}) + \frac{1}{N^+} \sum_{R^+} (\frac{1}{1 + e^{-\text{concat}(\mathcal{A}_1, \mathcal{A}_2)}}) \quad (14)$$

where concat means that attention maps are concatenated together, R^+ and R^- denote the set of positive and negative samples in the pseudo-affinity labels generated by refining pseudo-labels, while N^+ and N^- denote the number of R^+ and R^- respectively.

Total Loss: We adopt the cross-entropy loss as the segmentation loss \mathcal{L}_{seg} . Besides, to enhance the local consistency of segmentation results, we also adopt the commonly used regularization loss \mathcal{L}_{reg} (Tang et al. 2018; Zhang et al. 2021). The total loss is defined as:

$$\mathcal{L} = \lambda_1 (\mathcal{L}_{corr} + \mathcal{L}_{seg} + \mathcal{L}_{equ} + \mathcal{L}_{aux}) + \lambda_2 \mathcal{L}_{reg} + \lambda_3 \mathcal{L}_{cls} \quad (15)$$

where λ_1, λ_2 and λ_3 are empirically set to 0.1, 0.01 and 1 respectively.



Figure 3: Examples of segmentation results of AFA (Ru et al. 2022) and our method on PASCAL VOC validation images. The red box shows the difference.

Experiments

Experimental Settings

Dataset and evaluation metric: We evaluate our method on the commonly used PASCAL VOC 2012 (Everingham et al. 2010) dataset and MS COCO 2014 (Lin et al. 2014) dataset. The PASCAL VOC 2012 dataset consists of training, validation and test sets with a total of 21 semantic classes. Usually the SBD dataset (Hariharan et al. 2011) is used for the augmentation of the PASCAL VOC 2012 dataset, and the augmented dataset includes 10,582 images for training, 1,449 images for validation, and 1,464 images for testing. The MS COCO 2014 dataset (Lin et al. 2014) contains 81 classes with training set and validation set, each containing 82,081 and 40,137 images respectively. Note that we only use image-level labels for annotation. We adopt the mean Intersection-Over-Union (mIoU) as the evaluation metric.

Reproducibility: We implement our method using the PyTorch framework. For the training phase, we employ the AdamW optimizer with an initial learning rate set to 6×10^{-5} and a weight decay factor 0.01. We use simple data augmentation strategies such as random rescaling, random horizontal flipping, and random cropping. We set the batch size to 8 and the crop size to 512×512 . For VARM, we set (α, β) to (4, 0.01). Besides, we follow (Ru et al. 2022) and set the dilation rate of the dilated convolution to [1, 2, 4, 8, 12, 24]. The number of samples n for SCD loss is set to 40. For pseudo-label generation, we follow (Ru et al.

Method	Sup	Backbone	val	test
Fully-supervised methods				
DeepLab		R101	77.6	79.7
Segformer	<i>F</i>	MiT-B1	78.7	-
Multi-Stage weakly-supervised methods				
MCIS _{ECCV'2020}		R101	66.2	66.9
AuxSegNet _{ICCV'2021}	<i>I + S</i>	WR38	69.0	68.6
EPS _{CVPR'2021}		R101	70.9	70.8
SEAM _{CVPR'2020}		WR38	64.5	65.7
SC-CAM _{CVPR'2020}		R101	66.1	65.9
CDA _{ICCV'2021}	<i>I</i>	WR38	66.1	66.8
AdvCAM _{CVPR'2021}		R101	68.1	68.0
CPN _{ICCV'2021}		R101	67.8	68.5
RIB _{NeurIPS'2021}		R101	68.3	68.6
End-to-End weakly-supervised methods				
EM _{ICCV'2015}		VGG16	38.2	39.6
MIL _{CVPR'2015}		-	42.0	40.6
CRF-RNN _{CVPR'2017}		VGG16	52.8	53.7
RRM _{AAAI'2020}		WR38	62.6	62.9
RRM+ _{AAAI'2020}	<i>I</i>	MiT-B1	63.5	-
1Stage _{CVPR'2020}		WR38	62.7	64.3
AA&LR _{ACMMM'2021}		WR38	63.9	64.8
AFA _{CVPR'2022}		MiT-B1	63.8	-
AFA+CRF _{CVPR'2022}		MiT-B1	66.0	66.3
TSCD (Ours)		MiT-B1	65.0	65.2
TSCD + CRF (Ours)	<i>I</i>	MiT-B1	67.3	67.5

Table 1: Semantic segmentation comparisons of PASCAL VOC 2012 dataset. *F*: full supervision; *S*: saliency maps; *I*: image-level labels.

2022) and set the two background thresholds to 0.55 and 0.35, respectively. For the experiments on PASCAL VOC 2012, the total number of iterations is set to 20,000, with 2,000 iterations warmed up for the classification branch. For the experiments of MS COCO 2014, the total number of iterations is set to 80,000, with 5,000 iterations of warm-up.

Network Configuration: Considering the training efficiency, our transformer backbone adopts the recent Segformer (Xie et al. 2021). For the encoder, we adopt the Mix Transformer (MiT) (Xie et al. 2021) and initialize the parameters with ImageNet-1k pretrained weights. Then we use the MLP decoder header (Xie et al. 2021) as the decoder to predict the refined pseudo-labels.

Comparisons with State-of-the-Arts

PASCAL VOC 2012 dataset: For a fair comparison, we follow the evaluation code of (Ru et al. 2022). We report mIoU values of state-of-the-art semantic segmentation methods and our method’s mIoU values on VOC 2012 validation and test images in Table 1. It is observed that our TSCD outperforms all the other state-of-the-art end-to-end methods in VOC 2012. These end-to-end methods include 1Stage (Araslanov and Roth 2020) and AFA (Ru et al. 2022), etc. R101 means using ResNet101 (He et al. 2016) as the backbone, and WR38 means using WideResNet38 (Wu, Shen, and Van Den Hengel 2019). Our method also achieves




Figure 4: Examples of segmentation results of our method on COCO val images.

comparable performance to recent multi-stage methods, including SEAM (Wang et al. 2020), SC-CAM (Chang et al. 2020) and CDA (Su et al. 2021), etc. Compared to RRM (Zhang et al. 2020a) and AFA (Ru et al. 2022) using the same backbone, our method improves mIoU on the val set by 5% and 2%, respectively. Our TSCD even achieves 85.5% of its fully supervised counterpart, Segformer (Xie et al. 2021). This validates the effectiveness of our proposed Self Correspondence Distillation method (SCD) and Variation-aware Refine Module (VARM).

MS COCO 2014 dataset: The mIoU values for the semantic segmentation on the MS COCO 2014 dataset are presented in Table 2. We report the semantic segmentation performance of multi-stage methods (EPS (Lee et al. 2021b), AuxSegNet (Xu et al. 2021a), CDA (Su et al. 2021), CGNet (Kweon et al. 2021) and RIB (Lee et al. 2021a)) and end-to-end methods AFA (Ru et al. 2022). We can see that our TSCD significantly outperforms these state-of-the-art end-to-end methods and multi-stage methods except RIB (Lee et al. 2021a), reaching mIoU of 40.1%. Compared with AFA using the same backbone, the mIoU of our method is improved by 3%.

Ablation Studies

We conduct various ablation studies on the PASCAL VOC 2012 dataset.

Ablation studies of our TSCD framework: First, we verify the effectiveness of each part of our TSCD. Table 3 reports the quantitative results of ablation analysis on the PASCAL VOC 2012 validation set. Our transformer-based baseline model achieves an mIoU of 47.2%. Compared to the baseline, our VARM and SCD further significantly improve mIoU by 20% and 34%, respectively. Using the auxiliary loss \mathcal{L}_{aux} and equivariant regularization loss \mathcal{L}_{equ} brings the proposed framework to 64.2% mIoU and 65.0% mIoU,

Method	Sup	Backbone	mIoU(%)
Multi-Stage weakly-supervised methods			
AuxSegNet _{ICCV'2021}	I + S	WR38	33.9
EPS _{CVPR'2021}		R101	35.7
SEAM _{CVPR'2020}		WR38	31.9
CONTA _{NeurIPS'2020}		WR38	32.8
CDA _{ICCV'2021}	I	WR38	31.7
CGNet _{ICCV'2021}		WR38	36.4
RIB _{NeurIPS'2021}		R101	43.8
End-to-End weakly-supervised methods			
AFA _{CVPR'2022}	I	MiT-B1	38.0
AFA + CRF _{CVPR'2022}		MiT-B1	38.9
TSCD (Ours)	I	MiT-B1	39.2
TSCD + CRF (Ours)		MiT-B1	40.1

Table 2: Semantic segmentation comparisons of MS COCO 2014 validation images.

VARM	SCD	\mathcal{L}_{aux}	\mathcal{L}_{equ}	Method	val
				Baseline	47.2
✓					56.8
✓	✓			TSCD (Ours)	63.5
✓	✓	✓			64.2
✓	✓	✓	✓		65.0

Table 3: Ablation studies on our TSCD framework. Here ✓ indicates that this component is applied.

respectively.

Method	Backbone	val
R101-Base	R101	53.5
WR38-Base	WR38	56.2
R101-SCD	R101	56.5
WR38-SCD	WR38	59.1
TSCD <i>w/o</i> SCD	MiT-B1	63.9
TSCD (Ours)	MiT-B1	65.0

Table 4: Ablation studies on our SCD.

The experiments in Table 3 validate the effectiveness of individual modules in our framework. We also visualized the CAM of the baseline and the CAMs after applying our VARM and SCD respectively. As shown in Figure 6, the results show that VARM can further refine CAM, and SCD can enable the network to capture complete semantic information and improve the quality of pseudo-labels.

Ablation studies on SCD: To verify the generalization and effectiveness of our SCD, we also apply it to different mainstream backbones for weakly supervised semantic segmentation. We named the ResNet101 (He et al. 2016) backbone we implemented for the WSSS task as R101-Base, and the application of SCD on the R101-Base as R101-SCD. Implementations based on the WideResNet38 (Wu, Shen, and




Figure 5: Visualization of pseudo-label refinement via VARM. Pseudo-labels are generated using CAM and our baseline.

k^{rgb}	\mathcal{V}	Method	val
		Baseline	47.2
✓		PAMR (Araslanov and Roth 2020)	54.3
✓		PAR (Ru et al. 2022)	54.9
✓		VARM	55.0
✓	✓	VARM	56.8

Table 5: Ablation studies on VARM and comparison with other refinement methods.




Figure 6: Visualization examples of CAMs. From top to bottom: (a) CAMs generated by our baseline; (b) CAMs after applying our VARM; (c) CAMs after applying our SCD.

Van Den Hengel 2019) backbone are named WR38-Base and WR38-SCD, similar as ResNet101. The quantitative results in Table 4 show that both R101-SCD and WR38-SCD improve the network performance compared to the original backbone network, reaching mIoU of 56.5% and 59.1% on the VOC val, respectively. After removing the SCD from our TSCD, the network’s ability to obtain comprehensive semantic information is weakened, resulting in a considerable decrease in the network’s performance.

Ablation studies on VARM: Our VARM aims to overcome local inconsistencies by perceiving pixel-level variation, thereby refining the initial pseudo-labels. Table 5 shows the impact of each component in VARM on the segmentation results, and the comparison with other refinement methods such as PAMR (Araslanov and Roth 2020) and PAR (Ru et al. 2022). We use the CAM generated by the transformer-based baseline as the initial pseudo-label. The quantitative results in Table 5 show that our VARM can effectively improve the segmentation accuracy. Compared to baseline, our VARM improves mIoU by 20% (47.2 \rightarrow 56.8). Our VARM, PAMR (Araslanov and Roth 2020) and PAR (Ru et al. 2022) are all based on dilated pixel-adaptive convolutions, and our method outperforms PAMR (Araslanov and Roth 2020) and PAR (Ru et al. 2022), which validates the effectiveness of

VARM. In addition, Figure 5 shows that our VARM can enhance the local consistency of pseudo-labels and make the segmentation boundaries clearer, which makes the refined pseudo-labels closer to the ground truth. See supplementary material for more details for our SCD and VARM.

Conclusion

We address the challenges of WSSS using image-level class labels in obtaining comprehensive semantic information and high segmentation accuracy, and propose a simple and novel Self Correspondence Distillation method (SCD) to refine CAM without additional labels. Our SCD uses the network’s own CAM feature correspondence as the distillation target to enhance the feature learning process of the network. To further refine the initial pseudo-labels, we design the Variation-aware Refine Module (VARM), which enhances the local consistency of pseudo-labels by perceptual pixel-level variation. Based on SCD and VARM, an end-to-end Transformer-based framework (TSCD) for WSSS is provided. Extensive experiments demonstrate the effectiveness of our method, achieving the state-of-the-art performance on the PASCAL VOC 2012 and MS COCO 2014 datasets.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (Nos. U21A20515, 62271074, 61971418, U2003109, 62171321, 62071157 and 62162044) and in part by the Open Research Fund of Key Laboratory of Space Utilization, Chinese Academy of Sciences (No. LSU-KFJJ-2021-05), Open Research Projects of Zhejiang Lab (No. 2021KE0AB07), and this work was supported by the Open Projects Program of National Laboratory of Pattern Recognition.

References

- Ahn, J.; Cho, S.; and Kwak, S. 2019. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2209–2218.
- Araslanov, N.; and Roth, S. 2020. Single-stage semantic segmentation from image labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4253–4262.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Chang, Y.-T.; Wang, Q.; Hung, W.-C.; Piramuthu, R.; Tsai, Y.-H.; and Yang, M.-H. 2020. Weakly-supervised semantic segmentation via sub-category exploration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8991–9000.
- Chen, T.; Kornblith, S.; Swersky, K.; Norouzi, M.; and Hinton, G. E. 2020. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33: 22243–22255.
- Collins, E.; Achanta, R.; and Susstrunk, S. 2018. Deep feature factorization for concept discovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 336–352.
- Dong, J.; Cong, Y.; Sun, G.; Zhong, B.; and Xu, X. 2020. What Can Be Transferred: Unsupervised Domain Adaptation for Endoscopic Lesions Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4022–4031.
- Du, Y.; Fu, Z.; Liu, Q.; and Wang, Y. 2021. Weakly Supervised Semantic Segmentation by Pixel-to-Prototype Contrast. *arXiv preprint arXiv:2110.07110*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2): 303–338.
- Hamilton, M.; Zhang, Z.; Hariharan, B.; Snavely, N.; and Freeman, W. T. 2022. Unsupervised Semantic Segmentation by Distilling Feature Correspondences. *International Conference on Learning Representations*.
- Hariharan, B.; Arbeláez, P.; Bourdev, L.; Maji, S.; and Malik, J. 2011. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, 991–998. IEEE.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *stat*, 1050: 9.
- Hou, Q.; Jiang, P.; Wei, Y.; and Cheng, M.-M. 2018. Self-erasing network for integral object attention. *Advances in Neural Information Processing Systems*, 31.
- Hou, Y.; Ma, Z.; Liu, C.; and Loy, C. C. 2019. Learning lightweight lane detection cnns by self attention distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1013–1021.
- Kweon, H.; Yoon, S.-H.; Kim, H.; Park, D.; and Yoon, K.-J. 2021. Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6994–7003.
- Lee, J.; Choi, J.; Mok, J.; and Yoon, S. 2021a. Reducing Information Bottleneck for Weakly Supervised Semantic Segmentation. *Advances in Neural Information Processing Systems*, 34.
- Lee, J.; Kim, E.; Lee, S.; Lee, J.; and Yoon, S. 2019. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5267–5276.
- Lee, S.; Lee, M.; Lee, J.; and Shim, H. 2021b. Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5495–5505.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- O Pinheiro, P. O.; Almahairi, A.; Benmalek, R.; Golemo, F.; and Courville, A. C. 2020. Unsupervised learning of dense visual representations. *Advances in Neural Information Processing Systems*, 33: 4489–4500.
- Ru, L.; Zhan, Y.; Yu, B.; and Du, B. 2022. Learning Affinity from Attention: End-to-End Weakly-Supervised Semantic Segmentation with Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Su, H.; Jampani, V.; Sun, D.; Gallo, O.; Learned-Miller, E.; and Kautz, J. 2019. Pixel-adaptive convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11166–11175.
- Su, Y.; Sun, R.; Lin, G.; and Wu, Q. 2021. Context decoupling augmentation for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7004–7014.

- Tang, M.; Perazzi, F.; Djelouah, A.; Ben Ayed, I.; Schroers, C.; and Boykov, Y. 2018. On regularized losses for weakly-supervised cnn segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 507–522.
- Wang, C.; Xu, R.; Xu, S.; Meng, W.; Xiao, J.; Peng, Q.; and Zhang, X. 2022a. Softgan: Towards Accurate Lung Nodule Segmentation via Soft Mask Supervision. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Wang, C.; Xu, R.; Xu, S.; Meng, W.; and Zhang, X. 2022b. CNDesc: Cross Normalization for Local Descriptors Learning. *IEEE Transactions on Multimedia*.
- Wang, C.; Xu, R.; Xu, S.; Meng, W.; and Zhang, X. 2022c. DA-Net: Dual Branch Transformer and Adaptive Strip Upsampling for Retinal Vessels Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 528–538. Springer.
- Wang, C.; Xu, R.; Zhang, Y.; Xu, S.; Meng, W.; Fan, B.; and Zhang, X. 2022d. MTLDesc: Looking Wider to Describe Better. *arXiv preprint arXiv:2203.07003*.
- Wang, C.; Xu, R.; Zhang, Y.; Xu, S.; and Zhang, X. 2021. Retinal Vessel Segmentation Via Context Guide Attention Net With Joint Hard Sample Mining Strategy. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 1319–1323. IEEE.
- Wang, Y.; Zhang, J.; Kan, M.; Shan, S.; and Chen, X. 2020. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12275–12284.
- Wu, T.; Huang, J.; Gao, G.; Wei, X.; Wei, X.; Luo, X.; and Liu, C. H. 2021. Embedded discriminative attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16765–16774.
- Wu, Z.; Shen, C.; and Van Den Hengel, A. 2019. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90: 119–133.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34.
- Xu, L.; Ouyang, W.; Bennamoun, M.; Boussaid, F.; Sohel, F.; and Xu, D. 2021a. Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6984–6993.
- Xu, R.; Li, Y.; Wang, C.; Xu, S.; Meng, W.; and Zhang, X. 2022a. Instance segmentation of biological images using graph convolutional network. *Engineering Applications of Artificial Intelligence*, 110: 104739.
- Xu, R.; Wang, C.; Fan, B.; Zhang, Y.; Xu, S.; Meng, W.; and Zhang, X. 2022b. DOMAINDESC: Learning Local Descriptors With Domain Adaptation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2505–2509. IEEE.
- Xu, R.; Wang, C.; Xu, S.; Meng, W.; and Zhang, X. 2021b. DC-net: Dual context network for 2D medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 503–513. Springer.
- Xu, R.; Wang, C.; Zhang, J.; Xu, S.; Meng, W.; and Zhang, X. 2023. RSSFormer: Foreground Saliency Enhancement for Remote Sensing Land-Cover Segmentation. *IEEE Transactions on Image Processing*.
- Zhang, B.; Xiao, J.; Wei, Y.; Sun, M.; and Huang, K. 2020a. Reliability does matter: An end-to-end weakly supervised semantic segmentation approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12765–12772.
- Zhang, D.; Zhang, H.; Tang, J.; Hua, X.-S.; and Sun, Q. 2020b. Causal intervention for weakly-supervised semantic segmentation. *Advances in Neural Information Processing Systems*, 33: 655–666.
- Zhang, X.; Peng, Z.; Zhu, P.; Zhang, T.; Li, C.; Zhou, H.; and Jiao, L. 2021. Adaptive Affinity Loss and Erroneous Pseudo-Label Refinement for Weakly Supervised Semantic Segmentation. In *Proceedings of the 29th ACM International Conference on Multimedia*, 5463–5472.
- Zhao, M.; Yan, W.; Xu, R.; Zhi, D.; Jiang, R.; Jiang, T.; Calhoun, V. D.; and Sui, J. 2021. An attention-based hybrid deep learning framework integrating temporal coherence and dynamics for discriminating schizophrenia. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 118–121. IEEE.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.