



Knowledge Transfer with Simulated Inter-image Erasing for Weakly Supervised Semantic Segmentation

Tao Chen¹ , Xiruo Jiang¹ , Gensheng Pei¹ , Zeren Sun¹ ,
Yucheng Wang² , and Yazhou Yao¹

¹ Nanjing University of Science and Technology, Nanjing 210094, China
{taochen, jiangxiruo622, peigsh, zereens, yazhou.yao}@njjust.edu.cn

² Horizon Robotics, Beijing 100089, China
yucheng.wang@horizon.cc

Abstract. Though adversarial erasing has prevailed in weakly supervised semantic segmentation to help activate integral object regions, existing approaches still suffer from the dilemma of under-activation and over-expansion due to the difficulty in determining when to stop erasing. In this paper, we propose a **K**nowledge **T**ransfer with **S**imulated **I**nter-**I**mage **E**rasing (KTSE) approach for weakly supervised semantic segmentation to alleviate the above problem. In contrast to existing erasing-based methods that remove the discriminative part for more object discovery, we propose a simulated inter-image erasing scenario to weaken the original activation by introducing extra object information. Then, object knowledge is transferred from the anchor image to the consequent less activated localization map to strengthen network localization ability. Considering the adopted bidirectional alignment will also weaken the anchor image activation if appropriate constraints are missing, we propose a self-supervised regularization module to maintain the reliable activation in discriminative regions and improve the inter-class object boundary recognition for complex images with multiple categories of objects. In addition, we resort to intra-image erasing and propose a multi-granularity alignment module to gently enlarge the object activation to boost the object knowledge transfer. Extensive experiments and ablation studies on PASCAL VOC 2012 and COCO datasets demonstrate the superiority of our proposed approach. Codes and models are available at <https://nust-machine-intelligence-laboratory.github.io/project-KTSE>.

Keywords: Weakly Supervised Learning · Semantic Segmentation · Inter-Image Erasing

1 Introduction

With huge progress in the era of deep learning, semantic segmentation has been widely applied in fields like autonomous driving and image editing [5, 41].

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-72946-1_25.




Fig. 1. (a) Previous adversarial erasing-based approaches typically suffer from the over-expansion problem, which is hard to constrain. (b) Different from their information removal strategy, we propose to add extra object knowledge from a paired image to weaken the current object activation. The localization ability of the network is then enhanced by improving the consequent less activated attention map through learning from the object knowledge of the anchor branch. (C) Result comparison.

Since deep learning thus far is data-driven, the model training typically involves a large abundance of labeled images. However, collecting accurate pixel-level annotation for segmentation tasks is highly labor-intensive and time-consuming [3, 9, 16]. Therefore, as a promising direction to alleviate such annotation burden, weakly supervised learning has attracted the attention of many researchers. In this paper, we focus on the weakly supervised semantic segmentation (WSSS) under the supervision of image-level labels.

The recent advance of WSSS typically follows the three-step pipeline of: 1) transforming image tags into pixel-level coarse labels, 2) refining the pseudo labels, and 3) training the final segmentation model with the refined labels. Speaking of segmentation label generation, the technique of class activation map (CAM) [71] has been the *de facto* paradigm for object localization with image-level labels. Unfortunately, the naive CAM can only highlight the most discriminative area of objects and thus its small and sparse activation will lead to incomplete object mining. Since the first step of seed generation is the foundation of later processes, numerous works have been developed to expand CAM activation for high-quality pseudo labels. Among them, adversarial erasing is one of the mainstream approaches, aiming to discover new and complementary object regions by masking the currently detected area in an adversarial manner [60, 67, 69]. For example, the pioneering work of AE-PSL [60] performs the erasing multiple times to progressively discover increasingly activated areas. The recent work of AEFT [67] proposes to learn the concept of erasing with the triplet loss between the input image, erased image, and negatively sampled image. Though improvement is obtained, existing adversarial erasing-based methods still suffer from the difficulty in determining when to stop erasing: excessive removal will lead to over-expansion and insufficient erasure will result in under-activation.

In this paper, we propose a **K**nowledge **T**ransfer with **S**imulated **I**nter-**I**mage **E**rasing (KTSE) approach for WSSS to alleviate the problem mentioned above. Existing erasing-based methods typically mask the most discriminative object part to force the classifier into seeking other regions. Based on the observation that information removal leads to more activation (which might also incur the

over-expansion illustrated in Fig. 1(a)), we act in a diametrically opposite way. Specifically, as shown in Fig. 1(b), we propose to simulate an inter-image erasing (SIE) scenario, introducing extra object information by concatenating a paired image and rendering the anchor image as the erased one. We then strengthen the object localization ability of the network by transferring the object knowledge from the anchor CAM to the consequent less activated localization map. Intuitively, our knowledge transfer with simulated inter-image erasing will not cause the over-expansion problem if the object in the anchor branch is well-discovered.

However, since the above knowledge transfer is bidirectional, it will also weaken object mining in anchor branch when learning from the sparse activation of simulated branch. Therefore, as shown in the right of Fig. 2, we propose a self-supervised regularization module for the anchor CAM feature to maintain its reliable activation in discriminative regions. Specifically, we first generate pseudo labels from the CAM feature by locating confident foreground and background area, which are used to directly supervise the learning of CAM features. However, extracted labels of the foreground are quite noisy in complex images with multiple categories of objects due to the inter-class confusion. Therefore, we also design an inter-class loss to implicitly encourage activation consistency in complex images to improve recognition of the inter-class object boundary.

Though our proposed self-supervised regularization can effectively constrain and facilitate the knowledge transfer in our simulated inter-image erasing module, the improvement of the network’s object localization ability will be severely limited by the quality of anchor CAM feature which is usually under-activated. Therefore, we resort to the traditional intra-image erasing and propose a multi-granularity alignment (MGA) module to gently expand the object activation without introducing much background noise. Specifically, as shown in the bottom of Fig. 2, we first leverage the image-level global alignment to distill the soft object confidence from the anchor branch to the masked one for object activation enlargement. Then, we leverage a pixel-level local alignment to transfer the newly discovered object information back to the anchor branch. With the gentle and constrained activation enlargement, the proposed multi-granularity alignment module further boosts the performance of our simulated inter-image erasing. As can be seen in Fig. 2, our SIE and MGA are designed in a symmetric way that play the adversarial role of expanding and constraining the anchor image activation respectively for mining complete and compact object regions.

Our contributions can be summarized as follows:

- (1) We propose a knowledge transfer with simulated inter-image erasing approach for weakly supervised semantic segmentation to alleviate the over-expansion problem of existing adversarial erasing-based methods.
- (2) A self-supervised regularization module is proposed to constrain the knowledge transfer for maintaining reliable activation and improving recognition of the inter-class object boundary.
- (3) We propose a multi-granularity alignment module to gently enlarge the object activation via image-level global alignment and pixel-level local alignment, boosting the knowledge transfer with simulated inter-image erasing.

2 Related Work

2.1 Weakly Supervised Semantic Segmentation

Weakly supervised semantic segmentation (WSSS) is the task that aims to obtain a high-performance segmentation model by learning only from weak labels. Compared with bounding boxes [16, 25, 27, 53], scribbles [37, 56] and points [3], the image-level label [1, 4, 10, 12, 24, 34, 42, 50, 66, 73] is the cheapest weak supervision to collect, making it the most popular annotation format for WSSS. The recent advance of WSSS [1, 24, 34, 66, 73] typically relies on the class activation technique of CAM [71] to locate the target object, which helps transform the image-level label to pixel-level dense annotation for segmentation network training. Since the object area highlighted in CAM is usually small and sparse, enlarging the CAM activation becomes the focus of the WSSS task. For example, RDC [61] and DRS [28] propagate the object information to less discriminative regions via dilated convolution with varied rates and attention suppression of most discriminative areas, respectively. Popular contrastive learning [18, 48, 73] and self-supervised approaches [8, 51, 59] are also applied to facilitate object mining by improving representation learning. Moreover, cross-image information [21, 57, 72] is exploited to help the classifier discover more object patterns to obtain consistent and integral target regions. Besides attention expansion, Ahn and Kwak propose AffinityNet [2] for CAM refinement, which propagates semantic affinity via a random walk to recover delicate shapes of objects. IRN [1] is then proposed to exploit class boundary detection for discovering the entire instance areas with more accurate boundaries.

2.2 Erasing-Based Approaches

Considering the original CAM technique usually leads to suboptimal performance due to failure to localize all object parts, erasing-based approaches are commonly employed to expand the highlighted region [30, 52, 60, 67, 69, 70]. They mask areas in a training image to force the network to seek other relevant parts. While direct erasing methods like Hide-and-Seek [52] hide image patches randomly, adversarial erasing-based approaches mask the most discriminative regions with the attention information, demonstrating more promising activation expansion potential. The first adversarial erasing-based work for WSSS is proposed by Wei *et al.* [60], which progressively discovers new object discriminative regions by erasing the currently activated area. Such a heuristic erasing strategy then attracts the attention of many researchers. For example, ACoL [69] proposes adversarial complementary learning with two adversary classifiers. Recently, Kweon *et al.* propose to apply the adversarial erasing framework to exploit the potential of the pre-trained classifier [30] and learn the concept of erasing with the triplet loss [67]. While adversarial erasing approaches achieve great success in enlarging the CAM activation, they tend to suffer from the over-expansion problem due to difficulty in determining when to stop erasing. Unlike the information removal strategy adopted in these approaches, we propose a

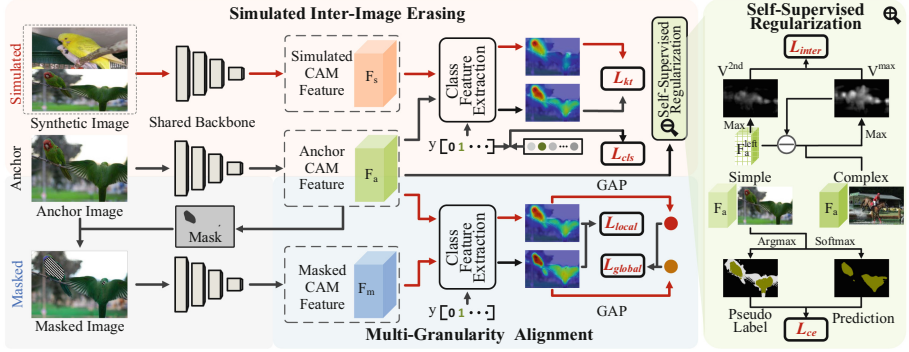


Fig. 2. The architecture of our proposed approach. We propose a simulated inter-image erasing (SIE) scenario where extra object information is introduced from another paired image. We then strengthen the object localization ability of the network by improving the consequent less activated localization map through learning object knowledge from the anchor image. A self-supervised regularization (SSR) module is also proposed to avoid weakening the anchor activation due to bidirectional alignment and improve the inter-class object boundary recognition for complex images. In addition, we propose a multi-granularity alignment (MGA) module to gently enlarge the object activation to further boost the object knowledge transfer.

simulated inter-image erasing scenario to weaken the activation by adding extra discriminative object information. The localization ability of network is then enhanced by improving the consequent less activated CAM through knowledge transfer.

3 Method

In this paper, we propose a knowledge transfer with simulated inter-image erasing approach for weakly supervised semantic segmentation to alleviate the over-expansion and under-activation problems of existing adversarial erasing-based methods. Our framework is illustrated in Fig. 2. As demonstrated in the anchor branch, we train a classification network with the given image-level weak labels, which consists of a backbone feature extractor and a pooling classification head. In contrast to existing adversarial erasing-based approaches that expand the activation through masking the most discriminative object part, as shown in the top-left of Fig. 2, we propose to simulate an inter-image erasing scenario where extra object information is introduced through concatenating a paired image. Object knowledge is then transferred from the anchor image to the consequent less activated localization map for strengthening the object localization ability of network. We also propose a self-supervised regularization module to avoid weakening the anchor activation due to bidirectional alignment and improve the inter-class object boundary recognition for complex multi-category images. In addition, we propose a multi-granularity alignment module to gently enlarge the object activation for boosting the object knowledge transfer.

3.1 CAM Generation

For the architecture of classification network, we follow the previous erasing-based work ACoL [69] and remove the final fully-connected layer. We set the output channel of the backbone to $C + 1$, where C is the number of foreground categories and 1 is added for the background. We can thus directly generate object localization maps from the class-aware CAM features F in the forward pass to facilitate the object knowledge transfer. To obtain CAM for each foreground class c , we feed the attention map F^c into a ReLU layer and then normalize it to the range from 0 to 1:

$$A^c = \frac{\text{ReLU}(F^c)}{\max(F^c)}. \quad (1)$$

To improve the quality of the CAM by capturing the global context and local details of the image, we adopt the Gated Pyramid Pooling (GPP) layer [67] as the final pooling head. With the classification logits q^c generated from the pooling head, we train the classification network with the multi-label soft margin loss as follows:

$$\mathcal{L}_{cls} = -\frac{1}{C} \sum_{c=1}^C y^c \log \sigma(q^c) + (1 - y^c) \log [1 - \sigma(q^c)]. \quad (2)$$

Here, $\sigma(\cdot)$ is the sigmoid function. y^c is the image-level label for the c -th class.

3.2 Simulated Inter-image Erasing

Due to the sparsity of the CAM activation, previous adversarial erasing-based methods typically enlarge the object activation by maintaining the classification confidence on masked images or features, where the most discriminative regions are erased. However, due to the absence of guidelines on when to stop erasing, these approaches easily incur over-expansion with excessive removal and might also still suffer from under-activation with insufficient erasure. Therefore, we propose a knowledge transfer with simulated inter-image erasing approach for weakly supervised semantic segmentation to alleviate the above problem.

In contrast to existing erasing-based approaches that remove the discriminative part for more object mining, we propose a simulated inter-image erasing (SIE) scenario to enable the network to benefit from introducing extra discriminative object information. Specifically, as shown in the top-left of Fig. 2, a larger synthetic image, created by concatenating the anchor image and another paired one, plays the role of the original image needing masking. Correspondingly, the anchor image can be treated as the erased one where the paired image is masked. Due to the introduction of extra discriminative object information, fewer object regions will be highlighted in the anchor part of the synthetic image. Therefore, we then strengthen the object localization ability of the network by transferring the more integral object knowledge from the original anchor CAM to the

consequently less activated anchor part of the synthetic image. The knowledge transfer loss for the simulated inter-image erasing can be formulated as follows:

$$\mathcal{L}_{kt} = \text{ReLU}(\hat{F}_a - \hat{F}_s), \quad \text{where } \hat{F}_a = \text{CFE}(F_a, y), \hat{F}_s = \text{CFE}(F_s, y). \quad (3)$$

Here, F_a and F_s denote the CAM features of the anchor and simulated branches, respectively. CFE represents class feature extraction, which extracts the feature channel related to the classes that exist in the image. The ReLU operation means that our alignment focuses on the more highlighted object region rather than the background. Note that no rotation or resize is performed for the concatenated image. We only introduce extra object information from the paired image and control other variables to guarantee the weakened activation for the anchor part. Intuitively, our knowledge transfer with simulated inter-image erasing will not cause the over-expansion problem when the object in the anchor branch is well-discovered. Our experiments demonstrate that such knowledge transfer can effectively strengthen the localization ability of the network to alleviate the under-activation, especially when multiple object instances co-exist in the image. Besides, benefiting from the bidirectional alignment, the anchor branch will also learn from the weakened activation of the simulated one and produce compact attention maps for small objects to mitigate the over-expansion problem.

3.3 Self-supervised Regularization

Our simulated inter-image erasing aims to improve the object localization ability of the network by maintaining the high activation given extra discriminative object information. However, the adopted knowledge transfer is bidirectional, which will also weaken the anchor branch object mining when learning from the sparse activation of the simulated branch. Since simply cutting the gradient propagation of the anchor branch makes the training unstable, we propose a self-supervised regularization module for the anchor CAM feature to maintain its reliable activation in discriminative regions. Specifically, with the generated CAM feature F_a and the corresponding CAM A_a , we first leverage two thresholds $\beta_h = 0.3$ and $\beta_l = 0.15$ to locate the confident foreground and background as follows (the subscript a is omitted for simplicity):

$$\hat{Y}_{i,j} = \begin{cases} \text{argmax}(F_{i,j,:}), & \text{if } \max(A_{i,j,:}) \geq \beta_h, \\ 0, & \text{if } \max(A_{i,j,:}) \leq \beta_l, \\ 255, & \text{otherwise,} \end{cases} \quad (4)$$

where 255 denote the ignored labels for uncertain regions and $\text{argmax}(\cdot)$ extracts the semantic class with the maximum activation value. After refining the pseudo label \hat{Y} with the pixel-adaptive refinement module [47], we leverage it to directly supervise the learning of CAM features with the cross-entropy loss (coordinates i, j are omitted):

$$\mathcal{L}_{ce} = - \sum_{c=0}^C \hat{Y}_{bg}^c \log \Gamma(F)^c + \omega_{fg} \hat{Y}_{fg}^c \log \Gamma(F)^c. \quad (5)$$

$\Gamma(\cdot)$ denotes the softmax function. We also leverage the pseudo label \hat{Y} generated from the anchor branch to guide the learning of the simulated CAM feature with the same cross-entropy loss. However, we notice that the extracted labels of the foreground are quite noisy, especially for complex images with multiple categories of objects where the inter-class boundary is blurred. Here, we define images with objects of only one category as simple images and others containing two or more categories as complex ones. Therefore, we discount the influence of foreground label \hat{Y}_{fg} with the weight of $\omega_{fg} = 0.0125$. Besides, as shown in the right of Fig. 2, we design an inter-class loss to implicitly encourage the activation consistency for the multiple foreground classes in complex images as follows:

$$\mathcal{L}_{inter} = \frac{\sum (V^{2nd} - V^{max}) \cdot M^f}{|M^f|}, \quad (6)$$

Here, M^f denotes the foreground mask obtained with a threshold $\beta_o = 0.2$ and $|\cdot|$ return the value of l_1 -norm. V^{max} and V^{2nd} denotes the largest and second largest activation value along the channel dimension of the CAM feature F_a , which can be obtained as follows:

$$V_{i,j}^{max} = \max(F_{i,j,:}), \quad F^{left} = \text{remove}(F, V_{max}), \quad V_{i,j}^{2nd} = \max(F_{i,j,:}^{left}). \quad (7)$$

Equation (6) improves the exclusivity of the class activation for each pixel, leading to more accurate boundaries between foreground objects. Our proposed self-supervised regularization module treats the simple and complex images differently, intending to maintain the reliable activation of the main branch without deteriorating the inter-class boundary caused by inaccurate pseudo labels.

3.4 Multi-granularity Alignment

Our proposed self-supervised regularization module can effectively constrain and boost the knowledge transfer in our simulated inter-image erasing module. However, the improvement of the network’s object localization ability will be severely limited by the quality of the anchor CAM, which is usually under-activated. Therefore, it is natural to resort to the traditional intra-image erasing, which aims to expand the object area. Following AEFT [67], we leverage a threshold of 0.6 to mask the most discriminative region. To avoid introducing too much background noise, we propose a multi-granularity alignment module to gently expand the object activation, which encourages the CAM expansion of the erased image and then transfers the learned object knowledge back to the anchor branch. As shown in the bottom of Fig. 2, we first input the anchor feature F_a and masked feature F_m into a class feature extraction module to obtain the corresponding features of the existing target categories (\hat{F}_a and \hat{F}_m). Following the experimental finding of AEFT [67] that rigid classification supervision [69] for the masked branch tends to trigger the over-expansion problem in adversarial erasing, we resort to the soft class confidence guidance from the anchor branch to the masked one. Specifically, we directly adopt a global average pooling (GAP)

operation to obtain the final class confidence for each branch and an image-level global alignment loss can be formulated as follows:

$$\mathcal{L}_{global} = GAP \left(ReLU \left(\hat{F}_a \right) \right) - GAP \left(ReLU \left(\hat{F}_m \right) \right),$$

where $\hat{F}_m = CFE(F_m, y)$. (8)

\hat{F}_a and CFE have been defined in Eq. (3). Another difference with rigid classification supervision is that benefiting from the class feature extraction, our global alignment loss only focuses on the logit confidence for classes that exist in the image, enabling more effective and efficient gradient propagation. Unlike AEFT [67] that utilizes the feature space of the GPP layer as embedding space for loss construction, our experiments reveal that direct alignment with the CAM feature leads to more promising performance. With class information learned from the anchor branch, we further leverage a pixel-level local activation alignment to transfer the newly discovered object information from the erased image to guide the anchor branch learning.

$$\mathcal{L}_{local} = ReLU \left(\hat{F}_m - \hat{F}_a \right). \quad (9)$$

Such pixel-level local alignment can also hinder the erased branch from activating the unwanted background area by learning from the lowly-activated regions of the anchor branch. With the gentle and constrained activation enlargement, the multi-granularity alignment module further boosts the performance of our proposed simulated inter-image erasing approach. As can be seen in Fig. 2, our proposed SIE and MGA branches are designed in a symmetric way that play the adversarial role of expanding and constraining the anchor image activation, respectively. Their synergically cooperation facilitates the convergence for mining complete and compact object regions.

3.5 Training Objective

The overall training loss is as follows:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{kt} + \mathcal{L}_{global} + \mathcal{L}_{local} + \mathcal{L}_{ce} + \lambda_{inter} \mathcal{L}_{inter}. \quad (10)$$

We empirically set $\lambda_{inter} = 0.005$ as the hyperparameter that controls the weight of inter-class loss.

4 Experiment

4.1 Datasets and Implementation Details

We evaluate our approach on the PASCAL VOC 2012 [19] and COCO [38] datasets. The PASCAL VOC 2012 dataset contains 21 classes (20 object categories and the background) for semantic segmentation. It has 10,582 images for

Table 1. Accuracy of pseudo-masks evaluated on PASCAL VOC 2012 training set.

Methods	Seed	w/IRN [1]
IRN [1] _{CVPR19}	48.3	66.3
MBMNet [40] _{MM20}	50.2	66.8
CONTA [68] _{NIPS20}	48.8	67.9
AdvCAM [32] _{CVPR21}	55.6	69.9
RIB [31] _{NIPS21}	56.6	70.6
ReCAM [14] _{CVPR22}	56.6	70.5
ESOL [35] _{NIPS22}	53.6	68.7
CLIMS [63] _{CVPR22}	56.6	70.5
AEFT [67] _{ECCV22}	56.0	71.0
ACR [15] _{CVPR23}	60.3	72.3
FPR [6] _{ICCV23}	63.8	-
KTSE (Ours)	67.0	73.8

Table 2. Quantitative comparisons on PASCAL VOC 2012 val and test sets with **VGG** backbone. Sup: Supervision, I: Image-level label, S: Saliency maps.

Methods	Sup	Val	Test
DRS [28] _{AAAI21}	I+S	63.6	64.4
GSM [36] _{AAAI21}	I+S	63.3	63.6
NSROM [66] _{CVPR21}	I+S	65.5	65.3
EPS [34] _{CVPR21}	I+S	67.0	67.3
L2G [24] _{CVPR22}	I+S	68.5	68.9
AffinityNet [2] _{CVPR18}	I	58.4	60.5
ICD [20] _{CVPR20}	I	61.2	60.9
BES [7] _{ECCV20}	I	60.1	61.1
ECS [55] _{ICCV21}	I	62.1	63.4
KTSE (Ours)	I	67.3	67.0

training (expanded with SBD [22]), 1,449 for validation and 1,456 for testing. The COCO dataset is a more challenging benchmark with 80 semantic classes and the background. Following previous works [36, 58, 68], we use the default train/val splits (80k images for training and 40k for validation) in the experiment. Mean intersection over union (mIoU) is adopted as the metric to evaluate the quality of our pseudo labels and segmentation results.

For the classification network, we follow the work of AEFT [67] and employ ResNet38 [62] as the backbone. We adopt a poly learning rate with an initial value of 10^{-2} and a power of 0.9. For the second stage training of WSSS, following the recent work of BECO [45], we adopt DeeplabV2 [5] as the segmentation network, which uses ResNet101 [23] as the backbone. The momentum and weight decay of the SGD optimizer are 0.9 and 10^{-4} . The initial learning rate is set to 10^{-2} and is decreased using polynomial decay. The segmentation model is trained for 80 epochs and 40 epochs on VOC and COCO datasets, respectively, with a common batch size of 16. We also follow the default setting of DeeplabV2 [5] and conduct experiments with VGG16 backbone for a more comprehensive comparison with previous approaches [7, 20, 24, 34, 55]. All our backbones are pre-trained on ImageNet [17].

4.2 Comparisons to the State-of-the-Arts

Accuracy of Pseudo-masks. We first report the quality of the segmentation seeds and the generated pseudo-masks derived from our approach. The comparison with other state-of-the-arts is presented in Table 1. As can be seen, our segmentation seed can reach the mIoU of 67.0%, bringing a gain of 18.7% compared to the baseline reported by IRN [1]. Our method can outperform the state-of-the-art method FPR [6] by 3.2%. With the further refinement of IRN [1], the mIoU of our generated pseudo-masks can arrive at 73.8%, surpassing the previous SOTA of AEFT [67] and ACR [15] by more than 1.5%.

Table 3. Quantitative comparisons on PASCAL VOC 2012 val and test sets with **ResNet** backbone. Sup: Supervision, I: Image-level label, S: Image-level label, S: Saliency maps, O: Saliency maps. Out-of-distribution data, L: Language.

Methods	Sup	Val	Test
DRS [28] _{AAAI21}	I+S	71.2	71.4
NSROM [66] _{CVPR21}	I+S	70.4	70.2
EPS [34] _{CVPR21}	I+S	71.0	71.8
AuxSeg [64] _{ICCV21}	I+S	69.0	68.6
PPC [18] _{CVPR22}	I+S	72.6	73.6
RCA [73] _{CVPR22}	I+S	72.2	72.8
L2G [24] _{CVPR22}	I+S	72.1	71.7
W-OoD [33] _{CVPR22}	I+O	69.8	69.9
CLIP-ES [39] _{CVPR23}	I+L	71.1	71.4
LPCAM [13] _{CVPR23}	I+S	71.8	72.1
AdvCAM [32] _{CVPR21}	I	68.1	68.0
CDA [54] _{ICCV21}	I	66.1	66.8
CSE [30] _{ICCV21}	I	68.4	68.2
RIB [31] _{NIPS21}	I	68.3	68.6
AMR [44] _{AAAI22}	I	68.8	69.1
MCT [65] _{CVPR22}	I	71.9	71.6
AFA [47] _{CVPR22}	I	66.0	66.3
SIPE [8] _{CVPR22}	I	68.8	69.7
ReCAM [14] _{CVPR22}	I	68.5	68.4
PPC [18] _{CVPR22}	I	67.7	67.4
ViT-PCM [46] _{ECCV22}	I	70.3	70.9
S-BCE [46] _{ECCV22}	I	70.0	71.3
AEFT [67] _{ECCV22}	I	70.9	71.7
TOCO [48] _{CVPR23}	I	69.8	70.5
OCR [15] _{CVPR23}	I	72.7	72.0
ACR [15] _{CVPR23}	I	71.9	71.9
BECO [45] _{CVPR23}	I	72.1	71.8
FPR [6] _{ICCV23}	I	70.3	70.1
MARS [26] _{ICCV23}	I	70.3	71.2
KTSE (Ours)	I	73.0	72.9

Methods	Sup	Val
IAL [58] _{IJCV20}	I+S	27.7
GWSM [36] _{AAAI21}	I+S	28.4
EPS [34] _{CVPR21}	I+S	35.7
I2CRC [73] _{TMM22}	I+S	31.2
RCA [73] _{CVPR22}	I+S	36.8
MDBA [11] _{TIP23}	I+S	37.8
BFBP [49] _{ECCV16}	I	20.4
SEC [29] _{ECCV16}	I	22.4
CONTA [68] _{NIPS20}	I	23.7
KTSE (Ours)	I	37.2

Table 5. Quantitative comparisons on COCO val set with **ResNet** backbone. Sup: Supervision, I: Image-level label, S: Saliency maps, L: Language.

Methods	Sup	Val
L2G [24] _{CVPR22}	I+S	44.2
CLIP-ES [39] _{CVPR23}	I+L	45.4
LPCAM [13] _{CVPR23}	I+S	42.1
OC-CSE [30] _{ICCV21}	I	36.4
CDA [54] _{ICCV21}	I	33.2
RIB [31] _{NIPS21}	I	43.8
MCT [65] _{CVPR22}	I	42.0
SIPE [8] _{CVPR22}	I	40.6
TOCO [48] _{CVPR23}	I	41.3
OCR [15] _{CVPR23}	I	42.5
ACR [15] _{CVPR23}	I	45.3
BECO [45] _{CVPR23}	I	45.1
FPR [6] _{ICCV23}	I	43.9
USAGE [43] _{ICCV23}	I	44.3
KTSE (Ours)	I	45.9

Accuracy of Segmentation Maps on PASCAL VOC 2012. Our segmentation results on the PASCAL VOC 2012 dataset with the backbone of VGG and ResNet are demonstrated in Table 2 and Table 3, respectively. As can be seen, with the VGG backbone, our approach achieves the performance of 67.3% on the validation set and 67.0% on the test set, better than other state-of-the-art methods with only image-level labels. Our segmentation results are also competitive to many approaches that rely on saliency maps, *e.g.*, NSROM [66] and EPS [34]. With the ResNet backbone, we can improve the results to 73.0% and 72.9% on

Table 6. Element-wise component analysis. SIE: Simulated Inter-Image Erasing, SSR: Self-Supervised Regularization, MGA: Multi-Granularity Alignment.

Base	SIE	SSR	MGA	mIoU
✓				54.2
✓	✓			57.1
✓		✓		57.4
✓			✓	60.2
✓	✓	✓		63.8
✓	✓	✓	✓	67.0

Table 7. Comparison of Our Multi-Granularity Alignment with Previous Erasing-based Approaches. GA: Global Alignment; LA: Local Alignment.

Method	mIoU
Base	54.2
Rigid Classification [69]	55.6
Soft GPP Feature [67]	54.9
Our GA	58.7
Our GA & LA	60.2

the validation and test sets, respectively. As can be seen in Table 3, our proposed approach can achieve better performance compared with recent SOTA methods. For example, our approach outperforms OCR [15] and ACR [15] by about 1% on the test set, demonstrating the superiority of our knowledge transfer with simulated inter-image erasing.

Accuracy of Segmentation Maps on COCO. For the more challenging COCO dataset, we provide performance comparisons with state-of-the-art WSSS methods using the backbone of VGG and ResNet in Table 4 and Table 5, respectively. As shown in Table 4, our proposed KTSE with VGG backbone can achieve the performance of 37.2% mIoU, much better than previous methods supervised with only image-level labels, *e.g.*, 13.5% mIoU higher than CONTA [68]. Besides, our results are also competitive compared to previous SOTA methods with additional saliency guidance like RCA [73] and MDBA [11]. With the ResNet backbone, our proposed KTSE reaches the best result of 45.9% mIoU compared to previous SOTA WSSS methods. Specifically, our approach outperforms ACR [15] and BECO [45] by 0.6% and 0.8% mIoU, respectively.

4.3 Ablation Studies

Element-Wise Component Analysis. In this part, we demonstrate the contribution of each component proposed in our approach to improving the quality of pseudo-masks. As mentioned in Sect. 3.1, we incorporate the GPP [67] into the classification network as our strong baseline. As shown in Table 6, with our simulated inter-image erasing (SIE) module, we can improve the accuracy of segmentation seeds from the baseline of 54.2% to 57.1%, which demonstrates the benefit of strengthening the localization ability of the network through weakening the activation with extra object information and then learning from the anchor branch. By maintaining reliable activation in discriminative regions, our self-supervised regularization (SSR) module can achieve the mIoU of 57.4%. Moreover, the combination of our proposed SIE and SSR can significantly boost

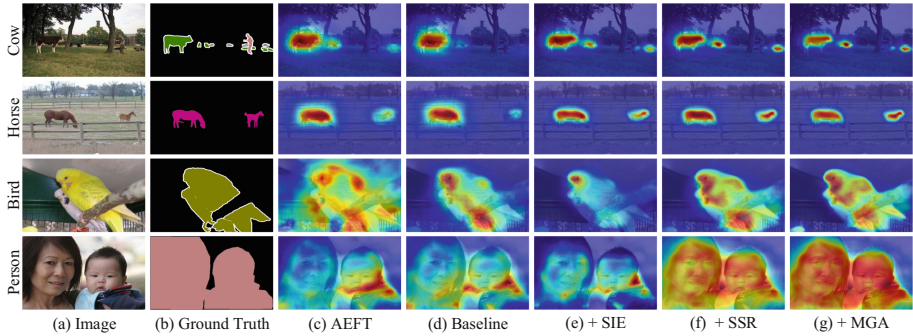


Fig. 3. Example localization maps on the PASCAL VOC 2012 training set. For each (a) image, we show (b) ground truth, localization maps produced by (c) previous work of AEFT [67], (d) our baseline, (e) baseline + SIE, (f) baseline + SIE + SSR, and (g) baseline + SIE + SSR + MGA. Best viewed in color.

the performance to 63.8% mIoU. With our proposed multi-granularity alignment module to gently expand the object activation, the mIoU of pseudo masks finally arrives at 67.0%, highlighting the importance of improving the quality of anchor CAM to boost the performance of our knowledge transfer with simulated inter-image erasing.

Some example localization maps on the PASCAL VOC 2012 training set can be viewed in Fig. 3. As seen from the first two rows, after strengthening the localization ability of the network with our simulated inter-image erasing (SIE) module, we can successfully expand the object activation and discover the less discriminative object in the image (*e.g.*, the smaller animals on the right). We can notice that our SIE can also help alleviate the over-expansion problem and generate more compact activation, *e.g.*, the cow in the first row. However, as demonstrated in the last two rows, the bidirectional alignment in SIE will also lead to more severe under-activation when the anchor branch learns from the sparse activation of the simulated one, especially for larger objects. Fortunately, our proposed self-supervised regularization (SSR) module can help maintain reliable activation in discriminative regions to guarantee effective knowledge transfer. Moreover, as seen in the last row, our multi-granularity alignment (MGA) module can further help discover the integral target region by gently expanding the object activation. Comparing the results of ours with that of AEFT [67], Fig. 3(g) vs Fig. 3(c), we can see that our proposed approach can significantly alleviate the under-activation and over-expansion of previous approaches due to insufficient or excessive erasure.

MGA vs. Previous Erasing-Based Approaches. For the multi-granularity alignment, we resort to the soft class confidence knowledge from the anchor branch to guide the gentle activation expansion of the masked image. In Table 7, we compare performance improvement from different guidance settings. As can be seen, our global alignment of CAM feature can improve the baseline from

54.2% to 58.7%, outperforming the rigid classification guidance in ACoL [69] and GPP feature alignment adopted in AEFT [67] by 3.1% and 3.8%, respectively. It highlights the importance of direct global alignment in CAM features rather than embedding space with extra projection layers. With our pixel-level local alignment, our multi-granularity alignment module can finally improve the performance to 60.2%, demonstrating the benefit of our gentle alignment strategy compared with previous erasing-based approaches.

Phenomenon of SIE. Based on the observation of classic adversarial erasing methods that removing the most discriminative region will lead to the extra activation of other object parts, the reversal thinking comes naturally—what if we introduce extra discriminative object information? The answer is the original highly activated area will become less discriminative and be lowly activated. We then enhance this weakened activation by learning from the original CAM. When the network learns to increase the attention for the less-discriminative region in the concatenated image, it will also learn to activate less-discriminative object region in the original image to locate more objects.

SIE vs. Data Augmentation. Though data augmentations like CutMix will also change the anchor image and lead to activation perturbation, they do not guarantee the weakened activation (which might cause over-expansion) for the anchor part of the concatenated image like ours. In contrast, we only introduce extra object information from the paired image and control other variables. Most importantly, our novelty lies in constructing the simulated inter-image erasing scenario (reversal thinking on existing masking methods) rather than simply augmenting the data input. It enhances the network’s localization ability by improving the consequent less activated attention map through learning from the object knowledge of the anchor branch. Accompanied by our proposed MGA, they are designed in a symmetric way that play the adversarial role of expanding and constraining the anchor image activation respectively for mining complete and compact object regions. Note that our approach significantly outperforms the CutMix-based augmentation method of CDA [54] with 73.0 *vs.* 66.1 on VOC (Table 3) and 45.9 *vs.* 33.2 on COCO (Table 5).

5 Conclusion

This paper proposed a Knowledge Transfer with Simulated Inter-Image Erasing (KTSE) approach for weakly supervised semantic segmentation. Specifically, in contrast to existing adversarial erasing-based methods that remove the discriminative part for mining less-discriminative areas, we proposed to simulate an inter-image erasing scenario where extra object information was added through concatenating a paired image. We then strengthened the object localization ability of the network by enhancing the consequent less activated localization map. In addition, we proposed a self-supervised regularization module to maintain reliable activation in the discriminative regions and improve the inter-class object boundary recognition for complex images. Moreover, we proposed a multi-granularity alignment module to gently enlarge the object activation via intra-image erasing for boosting the object knowledge transfer. Extensive experiments

and ablation studies were conducted on PASCAL VOC 2012 and COCO datasets to demonstrate the superiority of our proposed approach.

Acknowledgements. This work was supported by the National Natural Science Foundation of China (No. 62202227 and 62102182), Natural Science Foundation of Jiangsu Province (No. BK20220938 and BK20220934), China Postdoctoral Science Foundation (No. 2022M711635), Jiangsu Funding Program for Excellent Postdoctoral Talent (No. 2022ZB267), Fundamental Research Funds for the Central Universities (No. 30923010303).

References

1. Ahn, J., Cho, S., Kwak, S.: Weakly supervised learning of instance segmentation with inter-pixel relations. In: CVPR, pp. 2209–2218 (2019)
2. Ahn, J., Kwak, S.: Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In: CVPR, pp. 4981–4990 (2018)
3. Bearman, A., Russakovsky, O., Ferrari, V., Fei-Fei, L.: What’s the point: semantic segmentation with point supervision. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9911, pp. 549–565. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46478-7_34
4. Cai, X., Lai, Q., Wang, Y., Wang, W., Sun, Z., Yao, Y.: Poly kernel inception network for remote sensing detection. In: CVPR, pp. 27706–27716 (2024)
5. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.: DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. IEEE TPAMI **40**, 834–848 (2017)
6. Chen, L., Lei, C., Li, R., Li, S., Zhang, Z., Zhang, L.: FPR: false positive rectification for weakly supervised semantic segmentation. In: ICCV, pp. 1108–1118 (2023)
7. Chen, L., Wu, W., Fu, C., Han, X., Zhang, Y.: Weakly supervised semantic segmentation with boundary exploration. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12371, pp. 347–362. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58574-7_21
8. Chen, Q., Yang, L., Lai, J.H., Xie, X.: Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation. In: CVPR, pp. 4288–4298 (2022)
9. Chen, T., et al.: Semantically meaningful class prototype learning for one-shot image segmentation. IEEE TMM **24**, 968–980 (2021)
10. Chen, T., Yao, Y., Huang, X., Li, Z., Nie, L., Tang, J.: Spatial structure constraints for weakly supervised semantic segmentation. IEEE TIP **33**, 1136–1148 (2024)
11. Chen, T., Yao, Y., Tang, J.: Multi-granularity denoising and bidirectional alignment for weakly supervised semantic segmentation. IEEE TIP **32**, 2960–2971 (2023)
12. Chen, T., Yao, Y., Zhang, L., Wang, Q., Xie, G.S., Shen, F.: Saliency guided inter- and intra-class relation constraints for weakly supervised semantic segmentation. IEEE TMM **25**, 1727–1737 (2022)
13. Chen, Z., Sun, Q.: Extracting class activation maps from non-discriminative features as well. In: CVPR, pp. 3135–3144 (2023)

14. Chen, Z., Wang, T., Wu, X., Hua, X.S., Zhang, H., Sun, Q.: Class re-activation maps for weakly-supervised semantic segmentation. In: CVPR, pp. 969–978 (2022)
15. Cheng, Z., et al.: Out-of-candidate rectification for weakly supervised semantic segmentation. In: CVPR, pp. 23673–23684 (2023)
16. Dai, J., He, K., Sun, J.: BoxSup: exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: ICCV, pp. 1635–1643 (2015)
17. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: CVPR, pp. 248–255 (2009)
18. Du, Y., Fu, Z., Liu, Q., Wang, Y.: Weakly supervised semantic segmentation by pixel-to-prototype contrast. In: CVPR, pp. 4320–4329 (2022)
19. Everingham, M., Gool, L.V., Williams, C.K., Winn, J., Zisserman, A.: The PASCAL visual object classes (VOC) challenge. IJCV **88**, 303–338 (2010)
20. Fan, J., Zhang, Z., Song, C., Tan, T.: Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In: CVPR, pp. 4283–4292 (2020)
21. Fan, J., Zhang, Z., Tan, T., Song, C., Xiao, J.: CIAN: cross-image affinity net for weakly supervised semantic segmentation. In: AAAI, vol. 34, pp. 10762–10769 (2020)
22. Hariharan, B., Arbeláez, P., Bourdev, L.D., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: ICCV, pp. 991–998 (2011)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
24. Jiang, P.T., Yang, Y., Hou, Q., Wei, Y.: L2G: a simple local-to-global knowledge transfer framework for weakly supervised semantic segmentation. In: CVPR, pp. 16886–16896 (2022)
25. Jing, L., Chen, Y., Tian, Y.: Coarse-to-fine semantic segmentation from image-level labels. IEEE TIP **29**, 225–236 (2019)
26. Jo, S., Yu, I.J., Kim, K.: MARS: model-agnostic biased object removal without additional supervision for weakly-supervised semantic segmentation. arXiv preprint [arXiv:2304.09913](https://arxiv.org/abs/2304.09913) (2023)
27. Khoreva, A., Benenson, R., Hosang, J., Hein, M., Schiele, B.: Simple does it: weakly supervised instance and semantic segmentation. In: CVPR, pp. 876–885 (2017)
28. Kim, B., Han, S., Kim, J.: Discriminative region suppression for weakly-supervised semantic segmentation. In: AAAI, vol. 35, pp. 1754–1761 (2021)
29. Kolesnikov, A., Lampert, C.H.: Seed, expand and constrain: three principles for weakly-supervised image segmentation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9908, pp. 695–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_42
30. Kweon, H., Yoon, S.H., Kim, H., Park, D., Yoon, K.J.: Unlocking the potential of ordinary classifier: class-specific adversarial erasing framework for weakly supervised semantic segmentation. In: ICCV, pp. 6994–7003 (2021)
31. Lee, J., Choi, J., Mok, J., Yoon, S.: Reducing information bottleneck for weakly supervised semantic segmentation. In: NeurIPS, vol. 34, pp. 27408–27421 (2021)
32. Lee, J., Kim, E., Yoon, S.: Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In: CVPR, pp. 4071–4080 (2021)
33. Lee, J., Oh, S.J., Yun, S., Choe, J., Kim, E., Yoon, S.: Weakly supervised semantic segmentation using out-of-distribution data. In: CVPR, pp. 16897–16906 (2022)
34. Lee, S., Lee, M., Lee, J., Shim, H.: Railroad is not a train: saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In: CVPR, pp. 5495–5505 (2021)

35. Li, J., Jie, Z., Wang, X., Wei, X., Ma, L.: Expansion and shrinkage of localization for weakly-supervised semantic segmentation, vol. 35, pp. 16037–16051 (2022)
36. Li, X., Zhou, T., Li, J., Zhou, Y., Zhang, Z.: Group-wise semantic mining for weakly supervised semantic segmentation. In: AAAI, pp. 1984–1992 (2021)
37. Lin, D., Dai, J., Jia, J., He, K., Sun, J.: ScribbleSup: scribble-supervised convolutional networks for semantic segmentation. In: CVPR, pp. 3159–3167 (2016)
38. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
39. Lin, Y., et al.: CLIP is also an efficient segmenter: a text-driven approach for weakly supervised semantic segmentation. In: CVPR, pp. 15305–15314 (2023)
40. Liu, W., Zhang, C., Lin, G., Hung, T.Y., Miao, C.: Weakly supervised segmentation with maximum bipartite graph matching. In: ACM MM, pp. 2085–2094 (2020)
41. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR, pp. 3431–3440 (2015)
42. Pei, G., Chen, T., Jiang, X., Liu, H., Sun, Z., Yao, Y.: VideoMAC: video masked autoencoders meet convnets. In: CVPR, pp. 22733–22743 (2024)
43. Peng, Z., Wang, G., Xie, L., Jiang, D., Shen, W., Tian, Q.: USAGE: a unified seed area generation paradigm for weakly supervised semantic segmentation. arXiv preprint [arXiv:2303.07806](https://arxiv.org/abs/2303.07806) (2023)
44. Qin, J., Wu, J., Xiao, X., Li, L., Wang, X.: Activation modulation and recalibration scheme for weakly supervised semantic segmentation. In: AAAI, vol. 36, pp. 2117–2125 (2022)
45. Rong, S., Tu, B., Wang, Z., Li, J.: Boundary-enhanced co-training for weakly supervised semantic segmentation. In: CVPR, pp. 19574–19584 (2023)
46. Rossetti, S., Zappia, D., Sanzari, M., Schaerf, M., Pirri, F.: Max pooling with vision transformers reconciles class and shape in weakly supervised semantic segmentation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13690, pp. 446–463. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20056-4_26
47. Ru, L., Zhan, Y., Yu, B., Du, B.: Learning affinity from attention: end-to-end weakly-supervised semantic segmentation with transformers. In: CVPR, pp. 16846–16855 (2022)
48. Ru, L., Zheng, H., Zhan, Y., Du, B.: Token contrast for weakly-supervised semantic segmentation. In: CVPR, pp. 3093–3102 (2023)
49. Saleh, F., Aliakbarian, M.S., Salzmann, M., Petersson, L., Gould, S., Alvarez, J.M.: Built-in foreground/background prior for weakly-supervised semantic segmentation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9912, pp. 413–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_25
50. Sheng, M., Sun, Z., Cai, Z., Chen, T., Zhou, Y., Yao, Y.: Adaptive integration of partial label learning and negative learning for enhanced noisy label learning. In: AAAI, vol. 38, pp. 4820–4828 (2024)
51. Shimoda, W., Yanai, K.: Self-supervised difference detection for weakly-supervised semantic segmentation. In: ICCV, pp. 5208–5217 (2019)
52. Singh, K.K., Yu, H., Sarmasi, A., Pradeep, G., Lee, Y.J.: Hide-and-seek: a data augmentation technique for weakly-supervised localization and beyond. arXiv preprint [arXiv:1811.02545](https://arxiv.org/abs/1811.02545) (2018)
53. Song, C., Huang, Y., Ouyang, W., Wang, L.: Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In: CVPR, pp. 3136–3145 (2019)

54. Su, Y., Sun, R., Lin, G., Wu, Q.: Context decoupling augmentation for weakly supervised semantic segmentation. In: ICCV (2021)
55. Sun, K., Shi, H., Zhang, Z., Huang, Y.: ECS-Net: improving weakly supervised semantic segmentation by using connections between class activation maps. In: ICCV, pp. 7283–7292 (2021)
56. Vernaza, P., Chandraker, M.: Learning random-walk label propagation for weakly-supervised semantic segmentation. In: CVPR, pp. 7158–7166 (2017)
57. Wang, W., Sun, G., Van Gool, L.: Looking beyond single images for weakly supervised semantic segmentation learning. *IEEE TPAMI* **46**(3), 1635–1649 (2022)
58. Wang, X., Liu, S., Ma, H., Yang, M.H.: Weakly-supervised semantic segmentation by iterative affinity learning. *Int. J. Comput. Vis.* **128**, 1736–1749 (2020)
59. Wang, Y., Zhang, J., Kan, M., Shan, S., Chen, X.: Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In: CVPR, pp. 12275–12284 (2020)
60. Wei, Y., Feng, J., Liang, X., Cheng, M.M., Zhao, Y., Yan, S.: Object region mining with adversarial erasing: a simple classification to semantic segmentation approach. In: CVPR, pp. 1568–1576 (2017)
61. Wei, Y., Xiao, H., Shi, H., Jie, Z., Feng, J., Huang, T.: Revisiting dilated convolution: a simple approach for weakly-and semi-supervised semantic segmentation. In: CVPR, pp. 7268–7277 (2018)
62. Wu, Z., Shen, C., Van Den Hengel, A.: Wider or deeper: revisiting the ResNet model for visual recognition. *PR* **90**, 119–133 (2019)
63. Xie, J., Hou, X., Ye, K., Shen, L.: CLIMS: cross language image matching for weakly supervised semantic segmentation. In: CVPR, pp. 4483–4492 (2022)
64. Xu, L., Ouyang, W., Bennamoun, M., Boussaid, F., Sohel, F., Xu, D.: Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation. In: ICCV, pp. 6984–6993 (2021)
65. Xu, L., Ouyang, W., Bennamoun, M., Boussaid, F., Xu, D.: Multi-class token transformer for weakly supervised semantic segmentation. In: CVPR, pp. 4310–4319 (2022)
66. Yao, Y., et al.: Non-salient region object mining for weakly supervised semantic segmentation. In: CVPR, pp. 2623–2632 (2021)
67. Yoon, S.H., Kweon, H., Cho, J., Kim, S., Yoon, K.J.: Adversarial erasing framework via triplet with gated pyramid pooling layer for weakly supervised semantic segmentation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13689, pp. 326–344. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19818-2_19
68. Zhang, D., Zhang, H., Tang, J., Hua, X.S., Sun, Q.: Causal intervention for weakly-supervised semantic segmentation. In: NeurIPS, vol. 33 (2020)
69. Zhang, X., Wei, Y., Feng, J., Yang, Y., Huang, T.S.: Adversarial complementary learning for weakly supervised object localization. In: CVPR, pp. 1325–1334 (2018)
70. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: AAAI, vol. 34, pp. 13001–13008 (2020)
71. Zhou, B., Khosla, A., Lapedriza, À., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR, pp. 2921–2929 (2016)
72. Zhou, T., Li, L., Li, X., Feng, C.M., Li, J., Shao, L.: Group-wise learning for weakly supervised semantic segmentation. *IEEE TIP* **31**, 799–811 (2021)
73. Zhou, T., Zhang, M., Zhao, F., Li, J.: Regional semantic contrast and aggregation for weakly supervised semantic segmentation. In: CVPR, pp. 4299–4309 (2022)