Adaptive Affinity Loss and Erroneous Pseudo-Label Refinement for Weakly Supervised Semantic Segmentation

Submission ID: 2778

ABSTRACT

Semantic segmentation has been continuously investigated in the last ten years, and majority of the established technologies are based on supervised models. In recent years, image-level weakly supervised semantic segmentation (WSSS), including single- and multi-stage process, has attracted large attention due to data labeling efficiency. In this paper, we propose to embed affinity learning of multi-stage approaches in a single-stage model. To be specific, we introduce an adaptive affinity loss to thoroughly learn the local pairwise affinity. As such, a deep neural network is used to deliver comprehensive semantic information in the training phase, whilst improving the performance of the final prediction module. On the other hand, considering the existence of errors in the pseudo labels, we propose a novel label reassign loss to mitigate over-fitting. Extensive experiments are conducted on the PASCAL VOC 2012 dataset to evaluate the effectiveness of our proposed approach that outperforms other standard single-stage methods and achieves comparable performance against several multi-stage methods.

CCS CONCEPTS

• Computing methodologies \rightarrow Scene understanding.

KEYWORDS

weakly supervised, semantic segmentation, single-stage, adaptive affinity

ACM Reference Format:

Submission ID: 2778. 2021. Adaptive Affinity Loss and Erroneous Pseudo-Label Refinement for Weakly Supervised Semantic Segmentation . In *Chengdu* '21: ACM Multimedia, October 20–24, 2021, Chengdu. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Semantic segmentation aims to predict pixel-wise classification results on images, which is one of the vital tasks in computer vision and machine learning. With the development of deep learning, a variety of Convolutional Neural Network (CNN) based semantic segmentation methods [4, 5, 26] have achieved promising successes. However, these methods need to collect pixel-level labels, which are time consuming, laborious, and computationally expensive. To alleviate pixel-level annotations, many works have been devoted to weakly supervised semantic segmentation (WSSS) which applies

Chengdu '21, October 20–24, 2021, Chengdu

© 2021 Association for Computing Machinery

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00 https://doi.org/10.1145/1122445.1122456 weak supervisions by, e.g., image-level labels [1, 15, 19, 21, 30, 31, 42, 45, 46], scribbles [23, 39, 40], and bounding boxes [6, 18, 36]. In this paper, we focus on the most challenging case under weakly-supervised settings where image-level labels are utilized for training WSSS models.

Existing algorithms mainly consist of two ways to perform on image-level WSSS, by improving the performance of the initial Class Activation Map (CAM) [55], or refining the initial response with an additional model. Since CAM models lead to incomplete initial masks, some recent methods [3, 38, 44, 53] aim to obtain more precise and broader activation regions to deal with this problem. Different from that, early methods [15, 43, 47] follow the seedand-expand principle to refine the initial seeds (e.g. CAMs), while recent methods adopt an affinity learning method [1] for further enhancement.

Compared with the above methods, single-stage methods [2, 52] are dedicated to efficient training in a straightforward manner. In spite of substantial segmentation accuracy and streamlined process, these methods still fall behind most of the existing multi-stage models. In fact, the quality of the initial labels closely determines the final segmentation results. Figure 1 depicts certain diversity in initial labels: (*i*) Figure 1(a), there are positive initial results compared with the ground truth. (*ii*) Figure 1(b), mortal erroneous pixels exist. Most multi-stage methods severely rely on the fixed initial predictions that may limit their segmentation performance. As opposed to this, single-stage methods obtain initial seeds during training.

In this work, we leverage the merits of multi-stage methods to improve the performance of single-stage image-level semantic segmentation. Our intention is to refine the initial labels end-toend via *pixel-level affinity learning* and *erroneous pixel revision* to improve the segmentation results. To be precise, we first create the Adaptive Affinity (AA) loss to enhance the process of semantic propagation. Instead of using a fixed confidence factor (i.e., 1), we here adopt adaptive weights in pairwise connections, which further suppresses the excessive attention on 'fake' pixels. Second, we propose a novel Label Reassign (LR) loss that acts on the semantic propagation in the embedding space. This empowers the network to discern errors in initial labels, and thereby avoids over-fitting to the erroneous seeds.

Our proposed approach, stemming from the standard singlestage WSSS method [52], learns to improve the localization inference in a unified style. The main contributions of this work are summarized as follows:

• We present the AA loss to enhance the process of semantic propagation by computing connectivity while adaptively learning the local pairwise affinity.

• We propose the LR loss to identify erroneous initial predictions, which alleviates over-fitting on erroneous labels.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Chengdu '21, October 20-24, 2021, Chengdu



- -

Figure 1: Conceptual illustration of initial predictions. Different colors represent confident areas of object classes and background: red for airplane, orange for table, and black for background. The neutral region is color-coded in white. (a) Ideal Initial Prediction, (b) Semantic Ambiguity. The general refining methods often have better results on (a) while several methods misfitted marks in (b).

• Experiments on PASCAL VOC 2012 illustrate that our method achieves state-of-the-art performance in single-stage image-level WSSS.

2 RELATED WORK

2.1 Weakly-Supervised Semantic Segmentation

In this section, we discuss the algorithms for weakly-supervised semantic segmentation using image-level labels, including singleand multi-stage methods.

Single-stage methods. Compared with multi-stage models, singlestage methods [12, 29, 31, 32, 51] do not have complicated training process. Due to its simplicity, these methods barely achieve sufficient segmentation accuracy. Recently, RRM [52] and SSSS [2] tailor appropriate steps and gain accuracy similar to that of multi-stage methods. RRM [52] generates the initial seeds during the model training, combining the information of shallow features to optimize the segmentation network. SSSS [2] presents a segmentation-based network and a self-supervised training scheme to deal with the inherent defects of CAMs.

Multi-stage methods. Existing multi-stage approaches firstly generate pixel-level seeds from image-level labels by CAM (or grad-CAM), and then refine these seeds. Similarly, several approaches [1, 15, 19, 43, 47] refine the initial cue via expanding the region of attention. SEC [19] uses three principles, i.e., seed, expand, and constrain, to refine the initial predictions. To improve the network training, MCOF [43] adopts a top-down and bottom-up framework to alternatively expands the labeled regions and optimizes the segmentation model. Similarly, the DSRG approach [15] refines the initial localization maps by adopting a seeded region growing approach during the network training. The MDC method [47] expands the seeds via utilizing multi-branches of convolutional layers with different dilation rates.

On the other hand, CAMs are mainly trained in a classification task and tend to be activated by part of the object, resulting in less satisfactory prediction. Researchers are committed to improving the quality of the initial seeds. Adversarial erasing [14, 45] is a popular CAM extension method, which erases the most discriminative part of CAMs, guides the network to learn classification features from the other regions, and expands activation maps. Instead of using the erasing scheme, SEAM [44] presents a self-supervised equivariant attention mechanism to reduce the semantic gap between the segmentation and the classification. ScE [3] adopts sub-category classification tasks to cluster the foreground regions. More recently, MCIS [38] adopts semantic attention information across images to locate objects. To separate the foreground objects from the background, ICD [7] proposes an efficient intra-class discrimination approach.

In summary, initial labels play a pivotal role on the image-level WSSS. Using this concept, we apply the concept of multi-stage methods in single-stage image-level WSSS in order to further improve the segmentation accuracy.

2.2 Learning pixel-level affinities

The deploy of pairwise pixel affinity has a long history in image segmentation [34, 37], which has a variety of variants recently. For example, in fully-supervised segmentation, AAF [17] uses adaptive affinity fields to capture and match the relations between neighboring pixels in the label space. Similarly, AffinityNet [1] considers pixel-wise affinity to propagate local responses to the neighbouring regions, one of the popular methods for refining initial seeds. Moreover, [8, 9] explore cross-image relationships to collect complementary information during inference. However, errors often exist in the generated semantic labels, while the 'fake' samples consequently adopt the same confidence to reinforce themselves in an affinity learning scheme, leading to over-fitting.

2.3 Metric Learning

Metric learning methods [15, 35, 48] aim to optimize the transferable embeddings by learning a distance-based prediction rule over the embeddings. One advantage of using metric learning is to learn better representations of points and transfer knowledge as much as possible in the embedding space. For example, center loss [48] has been brought up to serve as an auxiliary loss for softmax loss to learn more discriminative features. On the other hand, triplet loss [15] encourages embeddings of data points with the same label to get closer than those with different identities. Several variants of this work have also been proposed, such as [25, 28, 41]. To reduce the time-consuming mining of hard triplets and tremendous data expansion, a novel loss named triplet-center loss [11] is put forward. Furthermore, [24] employs a metric learning network to transfer information from limited labeled samples to a large number of unlabeled samples in the embedding space.

3 PROPOSED APPROACH

Let $I = \{(I_i, Y_i)\}_n$ be the set of training data, where I_n is the n^{th} input image, and $Y_n \in \{0, 1\}^C$ is the associated ground-truth image label for C semantic categories. As shown in Figure 2, data (I_i, Y_i) are sampled from I for training the classification branch, After having fed CAM and pseudo labels to the segmentation branch, $F^s \in \mathbb{R}^{H \times W \times C_1}$ and $F^c \in \mathbb{R}^{H \times W \times C_2}$ are obtained, each with the spatial dimension of $H \times W$ and different channel dimensions. In this work, we adopt the RRM framework [52] with classification



Figure 2: The proposed framework: Given an image, first of all, we create pseudo masks through the classification branch and then use the masks as segmentation supervision. Afterwards, we reassign the labeled point embedding and also derive the LR loss. Furthermore, we train the segmentation prediction with a novel adaptive affinity learning method and traditional cross entropy loss. The entire framework is optimized end-to-end during training.

and cross entropy losses as our baseline. As mentioned before, we transplant the concept of multi-stage methods in a single-stage image-level WSSS to improve the segmentation performance. In our approach, there are four components in the objective function: (1) Adaptive affinity loss, (2) label reassign loss, (3) standard cross entropy loss and (4) standard classification loss. The former two reflect the quality of initial labels and the relationship between image pixels.

3.1 Adaptive Affinity Loss

To adopt the labeled samples for the generation of pseudo masks, some multi-stage approaches have used affinity learning so that semantic relationships can be transferred from the initial labels to the other regions. We adopt this concept in our single-stage image-level WSSS so as to deliver accurate semantic information.

Standard Affinity Loss Our implementation is derived from the idea of AAF [17] which adopts affinity learning for network predictions instead of adding another network. Following the practice of AAF [17], we take advantage of the Standard Affinity (SA) loss, combined with cross entropy loss and semantic label correlation. Consequently, the SA loss encourages the network to learn inter- and inner-class pixel relationships within the labeled regions (i.e., for the pixels where initial labels and unlabeled regions are ignored).

Following the setting of [1], depending on whether the two pixels belong to the same label or the region they appear, the set of pixel pairs \mathcal{P} are divided into three subsets, where \mathcal{P}_{fg}^+ is set up for the same object label in the foreground, \mathcal{P}_{fg}^- for the same label in the background, and \mathcal{P}^- refers to different label pairs. Moreover, similar to AAF [17], we combine multiple 3 * 3 kernels with varying dilation rates to learn pixel relationships over the all locations.

To obtain the SA loss, we expect to apply our loss function with the same kernel size. Consequently, the Single Kernel Standard Affinity (SKSA) loss is defined in three cases:

$$\mathcal{L}_{\rm fg}^{+} = \frac{1}{\left|\mathcal{P}_{\rm fg}^{+}\right|} \sum_{(i,j)\in\mathcal{P}_{\rm fg}^{+}} W_{ij} \tag{1}$$

$$\mathcal{L}_{bg}^{+} = \frac{1}{\left|\mathcal{P}_{bg}^{+}\right|} \sum_{(i,j) \in \mathcal{P}_{bg}^{+}} W_{ij} \tag{2}$$

$$\mathcal{L}^{-} = \frac{1}{|\mathcal{P}^{-}|} \sum_{(i,j) \in \mathcal{P}^{-}} \max\left(0, m - W_{ij}\right)$$
(3)

where W_{ij} is the Kullback-Leibler divergence between the prediction probabilities. *m* is the margin of the separating force.

$$W_{ij} = D_{KL} \left(F_i^c \| F_j^c \right) \tag{4}$$

Furthermore, the SKSA loss \mathcal{L}_{SKSA} can be obtained as:

$$\mathcal{L}_{\text{SKSA}} = \mathcal{L}_{\text{fg}}^{+} + \mathcal{L}_{\text{bg}}^{+} + 2\mathcal{L}^{-}$$
(5)

Finally, the SA loss is defined as:

$$\mathcal{L}_{\text{SA}} = \sum_{k \in K} \mathcal{L}_{\text{SKSA}}^k \tag{6}$$

where the SA loss is superposed with different dilation rates K.

However, compared with the fully supervised semantic segmentation, initial labels out of single-stage methods may be of errors and unlabeled samples may be added to the process. In addition, the initial labels may be changed during the network training for the single-stage methods. Hence, we here propose an adaptive learning method to measure the validity of the affinity connections.



Figure 3: The refined results of initial labels with different epochs. The first row: outcome of the baseline. The second row: outcome of our proposed method.



Figure 4: The major difference between our proposed adaptive affinity learning and standard affinity learning is that the former combines edge connectivity with the activated confidence score of the corresponding vertices.

It is a common sense that the regions within the object or background incline to draw higher confidence scores in the initial predictions, while the regions near the borders receive lower scores. As illustrated at the left hand side of Figure 4, the standard learning methods assign the same connection weight to all the labeled samples. However, semantic confounding occurs frequently nearby the boundaries due to the misleading labels, producing negative impacts on semantic communication. The network learns the incorrect semantic relationship and keeps reinforcing it, leading to worse outcomes.

To address this issue, we propose to measure the connectivity of affinity based on the pixel activated score in order to enforce affinity learning on the region which is more reliable (i.e., internal region). As a result, we use the Adaptive Affinity (AA) loss to cooperate with pixel-wise cross entropy loss to better learn the pixel relationships.

In particular, we hypothesize that one pixel with a lower activated score is more unlikely to be true due to unpredictable semantic confounding. Given a sample (V_i, V_j) with V_i (the confidence activated score of pixel *i*) and V_j (the confidence activated score pixel *j*), we describe the probability of connectivity using a simple yet effective maximization function, defined as:

$$\Omega_{ii} = \max\left(V_i, V_i\right),\tag{7}$$

where the weight of connection Ω_{ij} is defined as the higher confidence score of the two points to penalize the connections within the boundaries.

Similar to the SKSA loss, the Single Kernel Adaptive Affinity (SKAA) loss is also defined in three cases:

$$\mathcal{L}_{\rm fg}^{+} = \frac{1}{\left|\mathcal{P}_{\rm fg}^{+}\right|} \sum_{(i,j)\in\mathcal{P}_{\rm fg}^{+}} \Omega_{ij} * W_{ij} \tag{8}$$

$$\mathcal{L}_{bg}^{+} = \frac{1}{\left|\mathcal{P}_{bg}^{+}\right|} \sum_{(i,j) \in \mathcal{P}_{bg}^{+}} \Omega_{ij} * W_{ij} \tag{9}$$

$$\mathcal{L}^{-} = \frac{1}{|\mathcal{P}^{-}|} \sum_{(i,j)\in\mathcal{P}^{-}} \max\left(0, \left(\Omega_{ij} * m - W_{ij}\right)\right)$$
(10)

Here, the SKAA loss \mathcal{L}_{SKAA} is defined as:

$$\mathcal{L}_{\text{SKAA}} = \mathcal{L}_{\text{fg}}^{+} + \mathcal{L}_{\text{bg}}^{+} + 2\mathcal{L}^{-}$$
(11)

Finally, the Adpative Affinity (AA) loss based on different dilation rates K is written as:

$$\mathcal{L}_{AA} = \sum_{k \in K} \mathcal{L}_{SKAA}^k \tag{12}$$

3.2 Erroneous Pseudo-Label Refinement

The detail of the label reassign loss is illustrated in the top-right of Figure 2, which learns semantic similarity between each feature embedding and the centroid embedding of each class to mitigate over-fitting to erroneous labels. For simplicity, we take the tripletcenter loss [11] as our prototype for model learning. Adaptive Affinity Loss and Erroneous Pseudo-Label Refinement for Weakly Supervised Semantic Segmentation

Motivation. As reported in [39], applying only few but correct labels may leads to satisfactory segmentation results. Although most methods apply high confidence thresholds to acquiring the initial predictions or updated regions in the iterative process, erroneous labels still exists somewhere. Upon this occurrence, the network suffers from over-fitting in the training process. To alleviate the above problem, we propose a novel Label Reassign (LR) loss based on metric learning.

Centroid Calculation. In order to better enhance the prototype, we calculate the per-class centroids related to network predictions, making the optimization result more robust. The per-class centroid is the weighted mean of its labeled samples in the embedding space:

$$c_k = \frac{1}{\sum_{i=1}^N \beta_i} \sum_{x_i \in I_k} \beta_i * F_{x_i}^s, \tag{13}$$

where β_i is the confidence value of the labeled pixel embedding x_i , which belongs to the set of feature embedding l_k , and $F_{x_i}^s$ is the embedding vector. Here, the confidence score can be obtained through calculating the network prediction probability. Then, the label of each pixel can be assigned to the class with the highest similarity value.

Label Reassign Loss. To address the problem of class imbalance, we split the embedding points into two parts, E_{bg} and E_{fg} for the background and the foreground, respectively. Then, the metric-based loss is computed as follows:

$$\mathcal{L}_{-} = \frac{1}{|E_{\rm bg}|} \sum_{x_i \in E_{\rm bg}} \alpha_i * \sum_{j=1}^{N} \max\left(0, n + D_{x_i c_j} - D_{x_i c_i}\right)$$
(14)

$$\mathcal{L}_{+} = \frac{1}{\left|E_{\mathrm{fg}}\right|} \sum_{x_i \in E_{\mathrm{fg}}} \alpha_i * \sum_{j=1}^{N} \max\left(0, n + D_{x_i c_j} - D_{x_i c_i}\right)$$
(15)

Where *N* is the class number (including the background) in each image, which is greater than or equal to 2, D(*) represents the cosine similarity and α_i is the modulation parameter for feature embedding *i*.

$$\alpha_{i} = \left(1 - \frac{D_{x_{i}c_{i}} - D_{x_{i}c_{j}}}{D_{x_{i}c_{i}} + D_{x_{i}c_{j}}}\right)^{\gamma},\tag{16}$$

where $D_{x_ic_i}$ is the distance between the reassigned pixel embedding and its nearest centroid c_i , while c_j is the second nearest centroid to x_i . γ is a concentration parameter. When $\gamma = 0$, the LR loss is identical to the prototype, and as γ increases, the LR loss focuses on the samples with a closer distance between $D_{x_ic_i}$ and $D_{x_ic_j}$, more likely to be affected by the erroneous labels.

Finally, the metric-based LR loss is defined as:

$$\mathcal{L}_{\mathrm{LR}} = \mathcal{L}_{-} + \mathcal{L}_{+} \tag{17}$$

EXPERIMENTAL WORK

4.1 Implementation Details

4

PASCAL VOC 2012. PASCAL VOC 2012 segmentation dataset is of 21 class annotations, i.e., 20 foreground objects and the background. The official dataset has 1464 images for training, 1449 for validation and 1456 for testing. Following the common experimental protocol for semantic segmentation, we take additional annotations from SBD [10] to build an augmented training set with 10582 images.



Figure 5: Visualization results of adaptive pairwise affinity learning. (a) Raw images. (b) Initial labels. (c) Baseline with only CE and CLS losses. (d) Baseline + SA loss (w/o Adaptive). (e) Baseline + AA loss. (f) Ground truth. Red rectangles highlight the results of incorrect semantic propagation.

Noting that only image-level classification labels are available during the network training. Mean intersection over union (mIoU) is considered as the evaluation criterion.

Initial mask generation. Similar to [2], we remove the masks for the classes with classification confidence < 0.1. Following the common practice [1, 52], we generate the initial masks after having applied dCRF [20] during the training.

Backbone Framework & **Network.** In order to validate the network, we adopt the proposed losses on RRM [52] with default parameters as our backbone framework due to its superior performance on single-stage WSSS and encouraging training efficiency. Identical to [52], we also use a WideResNet38 backbone network [49] provided by [1].

Data augmentation. Similar to RRM [52], we use random rescaling (falling in the range of (0.7, 1.3) w.r.t. the original image area) and horizontal flipping. Finally, the images are cropped to the size of 321 * 321.

Training Time & Memory. Our method requires a slight increase (less than 4%) in training time compared to RRM [52] (a single-stage technique), which is **0.8s** per image. Meanwhile, we reduce a large quantity of training time in comparison with all the multi-stage methods, especially NSROM [50] (over **4s** per image), due to their complex processing operations. The system is implemented in PyTorch and trained on four NVIDIA 1080ti GPUs with 11 GB memory.

Training. The final training objective is to combine all the above mentioned objectives, i.e., $L_{total} = \mathcal{L}_{CLS} + \mathcal{L}_{CE} + \lambda_1 \mathcal{L}_{AA} + \lambda_2 \mathcal{L}_{LR}$. We empirically set $\lambda_1 = 0.1, \lambda_2 = 0.1$. For the adaptive affinity loss, margin *m* is set to 3.0 with the 4-8-16-24 kernel size. For the label reassign loss, threshold *n* is set to 1 and used in the last two epochs to prevent the occurrence of over-fitting. For the feature dimension, we use $C_1 = 512, C_2 = 21$. We train our model with the initial learning rate of 1e-4 and the weight decay of 5e-4.

Chengdu '21, October 20-24, 2021, Chengdu



Figure 6: Visualization results of the label reassign (upsampling from feature map). Red rectangles highlight the improved regions predicted by label reassign which avoids over-fitting on noisy labels.

Table 1: Ablation study on Pascal VOC dataset. We study the impact of individual losses. Additionally, same as [2], we adopt ground-truth image-level labels to remove masks of any false positive class predicted by our model with \checkmark in the "rfpc" column.

CLS	CE	AA	LR	rfpc	IoU (train)	IoU (val)
\checkmark	\checkmark				63.7	60.0
\checkmark	\checkmark	\checkmark			66.7	63.2
\checkmark	\checkmark	\checkmark	\checkmark		67.4	63.9
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	68.2	65.8

4.2 Ablation Study and Analysis

We analyze the importance of the proposed losses, as shown in Table 1. By adding the proposed losses, the baseline can be improved significantly. Since the CNNs are not responsible for the delivery of correct semantic information, the adaptive affinity loss can lead us to better understanding this problem (from 63.7% to 66.7% IoU in the training set and 60.0% to 63.2% IoU in validation set). Moreover, the label reassign loss mitigates over-fitting on the erroneous labels, further improving the network performance (from 66.7% to 67.4% IoU in training set and 63.2% to 63.9% IoU in validation set). Finally, following the practice of [2], we remove any false positive class prediction to achieve the best results. As shown in Table 2, compared with the standard affinity loss, by adding adaptive connectivity, our approach significantly improves the IoU by 0.7% in the training set and 0.9% in the validation set.

Table 2: IoU on Pascal VOC dataset. w.r.t. the adaptive pairwise affinity loss. We compare the quality of SA (w/o adaptive) and the AA loss with different kernel sizes.

Size	4	8	12	24	36	IoU (train)	IoU (val)
	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	65.3	61.6
SA	\checkmark	\checkmark	\checkmark	\checkmark		66.0	62.3
	\checkmark	\checkmark	\checkmark			65.9	61.9
	↓	\checkmark	\checkmark	\checkmark	\checkmark	65.6	62.0
AA	\checkmark	\checkmark	\checkmark	\checkmark		66.7	63.2
	\checkmark	\checkmark	\checkmark			66.2	62.4

Table 3: IoU on Pascal VOC dataset. w.r.t. the adaptive pairwise affinity loss. We compare different modeling functions using the kernel size 4-8-12-24. $plus(V_i, V_j)$ means $(V_i + V_j)/2$.

MF	$max\left(V_{i},V_{j}\right)$	$min\left(V_i,V_j ight)$	$plus(V_i, V_j)$
IoU (train)	66.7	66.1	66.4
IoU (val)	63.2	62.5	62.7

Table 4: IoU on Pascal VOC dataset. w.r.t. the label reassign loss. We compare the quality of using different concentration parameters.

Ŷ	0	0.5	1	2	5
IoU (train)	66.9	67.1	67.2	67.4	67.0
IoU (val)	63.5	63.6	63.7	63.9	63.4

Labeled Regions Refinement. As shown in Figure 3, the baseline has little improvement on initial labels and over-fitting at the end of the network training, leading to the degradation of segmentation performance. On the other hand, our proposed method maintains correct semantic information, assigns correct labels in the unlabeled regions, and further improves the segmentation accuracy.

Adaptive Affinity Loss. In order to show the benefit of the proposed AA loss, we illustrate a set of comparison plots. As shown in Figure 5 (b), the initial labels have a number of erroneous prediction points. First, in Figure 5 (c), if we merely fit these points, the erroneous semantics will be passed to the unlabeled regions. This makes the results worse. Next, in Figure 5 (d), with pairwise affinity learning, the proposed framework learns the additional semantic information to set up reliable communication with the unlabeled regions, but over-fits to the wrong labeled regions as well. Finally, in Figure 5 (e), with the introduction of weights, the network learns the correct semantic information while mitigating over-fitting, and generates the results close to the ground truth.

Adaptive Affinity Loss and Erroneous Pseudo-Label Refinement for Weakly Supervised Semantic Segmentation

Chengdu '21, October 20-24, 2021, Chengdu



(b) validation

Figure 7: Qualitative results of our single-stage method on the PASCAL VOC 2012 dataset. From top to bottom of (a) training and (b) validation: input images, our segmentation results, and ground-truths.

Table 5: Performance on the PASCAL VOC 2012 validation set l	by severa	l single-stag	ge methods.
--	-----------	---------------	-------------

single stage Method	bkg	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbk	person	plant	sheep	sofa	train	tv	mIoU
RRM [52]	87.9	75.9	31.7	78.3	54.6	62.2	80.5	73.7	71.2	30.5	67.4	40.9	71.8	66.2	70.3	72.6	49.0	70.7	38.4	62.7	58.4	62.6
SSSS [2]	88.7	70.4	35.1	75.7	51.9	65.8	71.9	64.2	81.1	30.8	73.3	28.1	81.6	69.1	62.6	74.8	48.6	71.0	40.1	68.5	64.3	62.7
Ours	88.4	76.3	33.8	79.9	34.2	68.2	75.8	74.8	82.0	31.8	68.7	47.4	79.1	68.5	71.4	80.0	50.3	76.5	43.0	55.5	58.5	63.9

Modeling Function. To better understand the contribution of the adaptive affinity loss, Table 3 shows some common modeling functions related to confidence activated scores. For function $min(V_i, V_j)$, it penalizes the connections with lower confidence scores, avoiding most error samples. However, it also reduces the connections within the internal regions, leading to a marginal increase in the system performance. Next, for function $plus(V_i, V_j)$, it shows more smoothly changed connection scores in different regions, resulting in better performance, while blurring some of the objects unexpectedly. Moreover, function $max(V_i, V_j)$ weakens the connections within the boundaries where the confidence scores

of the two points are lower than the average, in which most semantic confounding occurs frequently whilst keeping clear semantic properties in the images. The significantly improved performance shows that adaptive affinity learning plays a positive role on WSSS.

Concentration Parameter. The results of using the LR loss with different concentration parameters are shown in Table 4. The LR loss introduces one new hyperparameter γ , the concentration parameter, that controls the strength of the modulation factor. When $\gamma = 0$, our loss is identical to the triplet-center loss. As γ increases, the LR loss focuses on more vulnerable samples, which are more likely to be overfitted. With $\gamma = 2$, the LR loss improves the IoU by 0.7 % in both training and val sets.



Figure 8: T-SNE visualization of point embeddings in the 2D space. (a) Feature embedding w/o label reassign (b) Feature embedding with label reassign.

Label Reassign. In order to understand how our learning mechanism improves the identification capacity of the network, we visualize the corresponding predictions after label reassignment (i.e., reassign the labeled pixels in the embedding space). As mentioned before, the network filters out some undesired noisy points, which helps alleviating over-fitting for the 'fake' labeled samples. Figure 6 shows our results, compared to the initial prediction, the label reassign produces better segmentation results. Moreover, we project the feature embedding via T-SNE [27]. The projected point embeddings are shown in Figure 8. Without the label reassign loss, semantic confounder exists in certain classes, demonstrating the effectiveness of our proposed loss.

4.3 Semantic Segmentation Performance

The proposed method performs better than the standard SSSS [2] and the RRM [52] on the validation set. For the test set, we achieve system performance similar to that of some multi-stage methods.

In Table 5, we show detailed results for each category on the validation set. In Figure 7, we present some examples of the final semantic segmentation results, where our results are close to the ground truthed segmentation.

In Table 6, we provide performance comparisons with the other image-level WSSS methods. Our method outperforms the other state of the art single-stage WSSS methods with the mIoU scores of 63.9 and 64.8 on the PASCAL VOC 2012 validation and test sets, respectively, and achieves similar performance to the multi-stage WSSS methods.

5 CONCLUSIONS

In this work, we have presented two novel losses to improve the performance of a single-stage WSSS. On the one hand, we proposed the Adaptive Affinity loss to model the relationship between image pixels so that the network can carry out accurate semantic propagation. On the other hand, we proposed the Label Reassign loss to identify the semantic information of image pixels in pseudo labels to ease the fitting of false labels. Our method produced better

nous.

Table 6: Mean IoU on the Pascal VOC validation and test sets, *I* indicates only image-level supervision with additional data *A* and saliency detection *S*.

Method	Backbone	Superv.	val	test
Multi-stage				
SEC [19]	VGG16	Ι	50.7	51.1
STC [46]	VGG16	I+S	49.8	51.2
AffinityNet [1]	WideResNet38	Ι	61.7	63.7
MCOF [43]	ResNet101	I+S	60.3	61.2
DSRG [15]	ResNet101	I+S	61.4	63.2
CIAN [8]	ResNet101	Ι	64.1	64.7
SeeNet [14]	ResNet101	I+S	63.1	62.8
FickleNet [21]	ResNet101	S	64.9	65.3
OAA [16]	ResNet101	Ι	63.9	65.6
ScE [3]	ResNet101	Ι	66.1	65.9
SEAM [44]	ResNet101	Ι	64.5	65.7
MCIS [38]	ResNet101	Ι	66.2	66.9
ICD [7]	ResNet101	Ι	64.1	64.3
CONTA [54]	ResNet101	Ι	66.1	66.7
GWSM [22]	ResNet101	I+S	68.2	68.5
NSROM [50]	ResNet101	I+S	68.3	68.5
Single-stage				
EM [29]	VGG16	Ι	38.2	39.6
MIL-LSE [31]	overfeat[33]	Ι	42.0	40.6
CRF-RNN [32]	VGG16	Ι	52.8	53.7
TransferNet [12]	VGG16	Α	52.1	51.2
WebCrawl [13]	VGG16	Α	58.1	58.7
RRM [52]	WideResNet38	Ι	62.6	62.9
SSSS [2]	WideResNet38	Ι	62.7	64.3
Ours	WideResNet38	Ι	63.9	64.8

segmentation results than several existing single-stage methods and a range of multi-stage methods.

REFERENCES

- Jiwoon Ahn and Suha Kwak. 2018. Learning Pixel-level Semantic Affinity with Image-level Supervision for Weakly Supervised Semantic Segmentation. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- [2] Nikita Araslanov and Stefan Roth. 2020. Single-Stage Semantic Segmentation from Image Labels. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4253–4262.
- [3] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. 2020. Weakly-Supervised Semantic Segmentation via Sub-Category Exploration. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 8991–9000.
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2017), 834–848.
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In Proceedings of the European conference on computer vision (ECCV). 801–818.

Adaptive Affinity Loss and Erroneous Pseudo-Label Refinement for Weakly Supervised Semantic Segmentation

- [6] Jifeng Dai, Kaiming He, and Jian Sun. 2015. Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In Proceedings of the IEEE international conference on computer vision. 1635–1643.
- [7] Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. 2020. Learning Integral Objects With Intra-Class Discriminator for Weakly-Supervised Semantic Segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4283–4292.
- [8] Junsong Fan, Zhaoxiang Zhang, Tieniu Tan, Chunfeng Song, and Jun Xiao. 2020. Cian: Cross-image affinity net for weakly supervised semantic segmentation. In Proceedings of the AAAI Conference on Artificial Intelligence. 10762–10769.
- [9] Ruochen Fan, Qibin Hou, Ming-Ming Cheng, Gang Yu, Ralph R Martin, and Shi-Min Hu. 2018. Associating inter-image salient instances for weakly supervised semantic segmentation. In Proceedings of the European Conference on Computer Vision (ECCV). 367–383.
- [10] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. 2011. Semantic contours from inverse detectors. In 2011 International Conference on Computer Vision. IEEE, 991–998.
- [11] Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, and Xiang Bai. 2018. Tripletcenter loss for multi-view 3d object retrieval. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 1945–1954.
- [12] Seunghoon Hong, Junhyuk Oh, Honglak Lee, and Bohyung Han. 2016. Learning transferrable knowledge for semantic segmentation with deep convolutional neural network. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 3204–3212.
- [13] Seunghoon Hong, Donghun Yeo, Suha Kwak, Honglak Lee, and Bohyung Han. 2017. Weakly supervised semantic segmentation using web-crawled videos. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 7322–7330.
- [14] Qibin Hou, PengTao Jiang, Yunchao Wei, and Ming-Ming Cheng. 2018. Selferasing network for integral object attention. In Advances in Neural Information Processing Systems. 549–559.
- [15] Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. 2018. Weakly-supervised semantic segmentation network with deep seeded region growing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 7014–7023.
- [16] Peng-Tao Jiang, Qibin Hou, Yang Cao, Ming-Ming Cheng, Yunchao Wei, and Hong-Kai Xiong. 2019. Integral object mining via online attention accumulation. In Proceedings of the IEEE International Conference on Computer Vision. 2070–2079.
- [17] Tsung-Wei Ke, Jyh-Jing Hwang, Ziwei Liu, and Stella X Yu. 2018. Adaptive affinity fields for semantic segmentation. In Proceedings of the European Conference on Computer Vision (ECCV). 587–602.
- [18] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. 2017. Simple does it: Weakly supervised instance and semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition. 876–885.
- [19] Alexander Kolesnikov and Christoph H Lampert. 2016. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European conference on computer vision*. Springer, 695–711.
- [20] Philipp Krähenbühl and Vladlen Koltun. 2011. Efficient inference in fully connected crfs with gaussian edge potentials. In Advances in neural information processing systems. 109–117.
- [21] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. 2019. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In Proceedings of the IEEE conference on computer vision and pattern recognition. 5267–5276.
- [22] Xueyi Li, Tianfei Zhou, Jianwu Li, Yi Zhou, and Zhaoxiang Zhang. 2020. Group-Wise Semantic Mining for Weakly Supervised Semantic Segmentation. arXiv preprint arXiv:2012.05007 (2020).
- [23] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. 2016. Scribblesup: Scribblesupervised convolutional networks for semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 3159–3167.
- [24] Wanyu Lin, Zhaolin Gao, and Baochun Li. 2020. Shoestring: Graph-Based Semi-Supervised Classification With Severely Limited Labeled Data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4174–4182.
- [25] Hongye Liu, Yonghong Tian, Yaowei Yang, Lu Pang, and Tiejun Huang. 2016. Deep relative distance learning: Tell the difference between similar vehicles. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2167–2175.
- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition. 3431–3440.
- [27] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. Journal of machine learning research 9, Nov (2008), 2579–2605.
- [28] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. 2016. Deep metric learning via lifted structured feature embedding. In Proceedings of the IEEE conference on computer vision and pattern recognition. 4004–4012.
- [29] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. 2015. Weakly-and semi-supervised learning of a deep convolutional network for

semantic image segmentation. In Proceedings of the IEEE international conference on computer vision. 1742–1750.

- [30] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. 2015. Constrained convolutional neural networks for weakly supervised segmentation. In Proceedings of the IEEE international conference on computer vision. 1796–1804.
- [31] Pedro O Pinheiro and Ronan Collobert. 2015. From image-level to pixel-level labeling with convolutional networks. In Proceedings of the IEEE conference on computer vision and pattern recognition. 1713–1721.
- [32] Anirban Roy and Sinisa Todorovic. 2017. Combining bottom-up, top-down, and smoothness cues for weakly supervised image segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 3529–3538.
- [33] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. 2013. Overfeat: Integrated recognition, localization and detection using convolutional networks. arXiv preprint arXiv:1312.6229 (2013).
- [34] Jianbo Shi and Jitendra Malik. 2000. Normalized cuts and image segmentation. IEEE Transactions on pattern analysis and machine intelligence 22, 8 (2000), 888– 905.
- [35] Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. In Advances in neural information processing systems. 1857–1865.
- [36] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. 2019. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 3136–3145.
- [37] X Yu Stella and Jianbo Shi. 2003. Multiclass spectral clustering. In null. IEEE, 313.
- [38] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. 2020. Mining crossimage semantics for weakly supervised semantic segmentation. arXiv preprint arXiv:2007.01947 (2020).
- [39] Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, and Christopher Schroers. 2018. Normalized cut loss for weakly-supervised cnn segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 1818–1827.
- [40] Paul Vernaza and Manmohan Chandraker. 2017. Learning random-walk label propagation for weakly-supervised semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition. 7158–7166.
- [41] Jian Wang, Feng Zhou, Shilei Wen, Xiao Liu, and Yuanqing Lin. 2017. Deep metric learning with angular loss. In Proceedings of the IEEE International Conference on Computer Vision. 2593–2601.
- [42] Xiang Wang, Sifei Liu, Huimin Ma, and Ming-Hsuan Yang. 2020. Weakly-Supervised Semantic Segmentation by Iterative Affinity Learning. International Journal of Computer Vision (2020), 1–14.
- [43] Xiang Wang, Shaodi You, Xi Li, and Huimin Ma. 2018. Weakly-supervised semantic segmentation by iteratively mining common object features. In Proceedings of the IEEE conference on computer vision and pattern recognition. 1354–1362.
- [44] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. 2020. Selfsupervised Equivariant Attention Mechanism for Weakly Supervised Semantic Segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 12275–12284.
- [45] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. 2017. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proceedings of the IEEE* conference on computer vision and pattern recognition. 1568–1576.
- [46] Yunchao Wei, Xiaodan Liang, Yunpeng Chen, Xiaohui Shen, Ming-Ming Cheng, Jiashi Feng, Yao Zhao, and Shuicheng Yan. 2016. Stc: A simple to complex framework for weakly-supervised semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence* 39, 11 (2016), 2314–2320.
- [47] Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. 2018. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 7268–7277.
- [48] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. 2016. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*. Springer, 499–515.
- [49] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. 2019. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition* 90 (2019), 119–133.
- [50] Yazhou Yao, Tao Chen, Guosen Xie, Chuanyi Zhang, Fumin Shen, Qi Wu, Zhenmin Tang, and Jian Zhang. 2021. Non-Salient Region Object Mining for Weakly Supervised Semantic Segmentation. arXiv preprint arXiv:2103.14581 (2021).
- [51] Yu Zeng, Yunzhi Zhuge, Huchuan Lu, and Lihe Zhang. 2019. Joint learning of saliency detection and weakly supervised semantic segmentation. In Proceedings of the IEEE International Conference on Computer Vision. 7223–7233.
- [52] Bingfeng Zhang, Jimin Xiao, Yunchao Wei, Mingjie Sun, and Kaizhu Huang. 2019. Reliability Does Matter: An End-to-End Weakly Supervised Semantic Segmentation Approach. arXiv preprint arXiv:1911.08039 (2019).
- [53] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. 2020. Causal intervention for weakly-supervised semantic segmentation. arXiv preprint arXiv:2009.12547 (2020).

Chengdu '21, October 20-24, 2021, Chengdu

- [54] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. 2020. Causal intervention for weakly-supervised semantic segmentation. arXiv preprint arXiv:2009.12547 (2020).
- [55] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2016. Learning deep features for discriminative localization. In Proceedings of the IEEE conference on computer vision and pattern recognition. 2921–2929.