

Seminar Learning for Click-Level Weakly Supervised Semantic Segmentation

Hongjun Chen¹, Jinbao Wang¹, Hong Cai Chen¹, Xiantong Zhen²,
Feng Zheng^{1*}, Rongrong Ji³, Ling Shao⁴

¹ Southern University of Science and Technology ² University of Amsterdam

³ Xiamen University ⁴ Inception Institute of Artificial Intelligence

Abstract

Annotation burden has become one of the biggest barriers to semantic segmentation. Approaches based on click-level annotations have therefore attracted increasing attention due to their superior trade-off between supervision and annotation cost. In this paper, we propose seminar learning, a new learning paradigm for semantic segmentation with click-level supervision. The fundamental rationale of seminar learning is to leverage the knowledge from different networks to compensate for insufficient information provided in click-level annotations. Mimicking a seminar, our seminar learning involves a teacher-student and a student-student module, where a student can learn from both skillful teachers and other students. The teacher-student module uses a teacher network based on the exponential moving average to guide the training of the student network. In the student-student module, heterogeneous pseudo-labels are proposed to bridge the transfer of knowledge among students to enhance each other's performance. Experimental results demonstrate the effectiveness of seminar learning, which achieves the new state-of-the-art performance of 72.51% (mIOU), surpassing previous methods by a large margin of up to 16.88% on the Pascal VOC 2012 dataset.

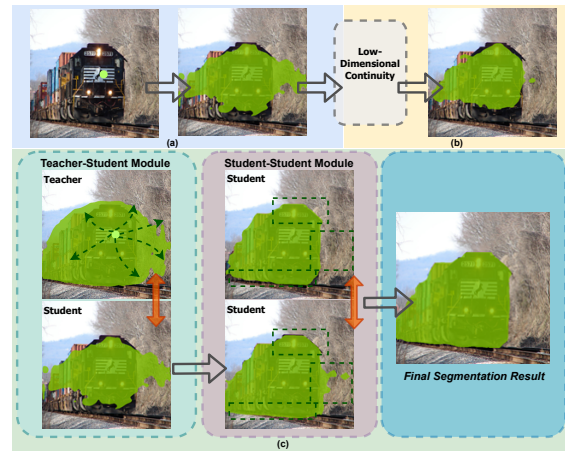


Figure 1. Weakly supervised segmentation with the click-level annotations. (a) A generic model trained only with click-level annotations overfits to the labels and cannot recognize the whole object. (b) Previous works (e.g. regularized loss) apply low-dimensional continuity information to the training, also failing to correctly segment the object. (c) With seminar learning, our teacher-student module enables the network to generalize to the whole object, as indicated by the arrows. Meanwhile, by integrating diverse information from the two networks, the student-student module can smooth the boundary area in the marked boxes.

1. Introduction

Semantic segmentation is a fundamental task, where each pixel of an image is labeled into a predefined set of classes. In the field of computer vision, it has made great progress in many applications, such as automatic driving, scene understanding, and medical diagnosis [30] [48]. Recently, deep convolutional neural networks (CNNs) have achieved remarkable success in a variety of semantic segmentation tasks [9, 30]. However, they require large amounts of pixel-level annotations for training. The acquisition process of pixel-level annotations is extremely time-

consuming and labor-intensive.

In order to alleviate the burden of annotations, weakly supervised semantic segmentation has become increasingly popular, as it only requires coarse annotations, such as box-level [12], image-level [33], scribble-level [24], or click-level [4] supervision. Among these, click-level supervision only annotates one pixel for each object in an image. Further, it not only provides valuable location information, but is also one of the cheapest form of weakly supervision [4]. It has high research potential in terms of the trade-off between information and time costs.

As commonly known, it is challenging to achieve satisfactory performance with limited supervision information during model training. For instance, with click-level pat-

*Corresponding author: Feng Zheng (Email: f.zheng@ieee.org). This work is supported by the National Natural Science Foundation of China under Grant No. 61972188.

terns, if only learning from one labeled pixel, the model cannot infer the entire range of an object, especially the edges, which will eventually weaken the segmentation performance. An effective way to compensate weak supervision information is to introduce more prior information. For example, ‘What’s the point’ [4] incorporates an objectness prior into network training, which helps distinguish between foreground and background. ‘ScribbleSup’ [24] uses an additional graphical model to propagate information from click-level annotations. ‘Regularized Loss’ [44] designs a regularization item based on dense conditional random field (CRF) for classifying nearby pixels with similar colors into the same category. These models only focus on low-dimensional continuity between labeled pixels and others pixels, which is limited to local annotation information in click-level supervision. Therefore, these models cannot properly segment the entire object and still underperform.

Considering the nature of click-level supervised semantic segmentation, we make two observations: 1) A large number of unlabeled pixels are not well used, but could provide broader information, which can expand the learning range of networks from a single annotated pixel to an entire object. 2) If a network is trained under different conditions, such as using different random seeds, the predictions will vary greatly. This uncertainty causes that different networks capture distinctive and diverse information, which could be aggregated to complement each other.

Inspired by these observations, we propose *seminar learning*, a novel learning paradigm for click-level weakly supervised semantic segmentation by introducing more effective information. The essence of our seminar learning is to complement the deficiency of networks by leveraging the knowledge provided from the predictions of other networks. As shown in Fig. 1, seminar learning framework consists of two components: teacher-student module and student-student module. Notably, the teacher-student module is exploited to expand the learning range of networks. We use an exponential moving average (EMA) based teacher network for generalized prediction and prevent the student network from overfitting to click-level labels, which has a similar workflow to semi-supervised mean-teacher [45] method. However, compared to mean-teacher, our module is able to operate on unlabeled pixels in each image instead of unlabeled images. The student-student module is applied to refine segmentation boundaries by aggregating diversity information of student networks. To improve the efficiency of information transfer, we propose heterogeneous pseudo-labels as bridges between student networks, which based on the prediction of a fully trained student to guide the other. In summary, we make several major contributions as follows:

- We propose a novel learning paradigm, called seminar learning, that can learn to leverage more supervisory information provided by a group of networks.

- We treat the click-level supervised semantic segmentation task as a semi-supervised pixel classification task per image, and propose a novel pixel consistency loss, which enables a student to learn from a teacher using unlabeled pixels.
- The novel concept of heterogeneous pseudo-labels is proposed, which is a more effective medium to enable the supervisory information to be shared among diverse networks by the student-student module.
- We conduct extensive experiments to verify the effectiveness of the proposed seminar learning, which outperforms previous SOTA works [44] by a large margin (from 55.63% to 72.51% in terms of the mIOU metric).

2. Related Work

Semantic segmentation. Semantic segmentation can be viewed as pixel-wise classification, where each pixel is assigned to a category. Using deep learning in semantic segmentation began when the fully convolutional network (FCN) [27] first arose. Most popular architectures are based on encoder-decoder models like FCN, such as SegNet [3], U-Net [37], MobileNet [39] and RefineNet [25]. Recent works have made great efforts to increase the receptive field of the network. For example, Chen *et al.* [10] extract long-range information without down-scaling the image by atrous convolutions. DeepLabv2 [9] introduces an atrous spatial pyramid pooling module. In this paper, we employ DeepLab-v3+ [10] as our backbone since it integrates various effective modules and outperforms previous models.

Semi-supervised learning. In general, semi-supervised learning tackles the training problem with only a small amount of labeled data. How to use plenty of unlabeled data becomes crucial for semi-supervised learning. A effective method is pseudo-labeling [22], which use the prediction generated by itself to supervise unlabeled data. Recently, there are many semi-supervised methods based on the conception of predictions consistent with perturbations. The Π -model [21] produces two noisy models, named the student model and the teacher model, and applies a consistency loss on the teacher and student predictions to adapt the model to the noise. Virtual adversarial training [31] is similar to the Π -model but it uses adversarial perturbation as noise. Mean-teacher [45] improves the teacher-student architecture. It obtains the teacher model by moving the average of the student model weights. This practice has also been followed by subsequent semi-supervised works [6, 29, 46, 40]. Besides, many algorithms utilize these semi-supervised methods for the image segmentation problem. For instance, [11] directly applies mean-teacher to the unlabeled data in semi-supervised semantic segmentation, while

The training procedure can be described as follows:

Training the ancillary model. The ancillary model is constructed by the teacher-student module. In this module, we only need to train the student network. The teacher network is obtained by the exponential moving average (EMA) of the student network. At the training iteration t , the EMA process is defined as

$$\theta'_t = \begin{cases} (1 - \frac{1}{t}) \times \theta'_{t-1} + \frac{1}{t} \times \theta_t, & 1 - \frac{1}{t} < \alpha \\ \alpha \theta'_{t-1} + (1 - \alpha) \theta_t, & otherwise, \end{cases} \quad (1)$$

where α is a smoothing coefficient hyperparameter, and θ' and θ are the weight of the teacher and student, respectively. To renew the weight of the teacher model quickly during the initial training iterations, we use absolute average instead of EMA when $1 - \frac{1}{t} < \alpha$.

The networks of the ancillary student and the ancillary teacher are randomly initialized with the same random seed. In each iteration of the training, we input training images to both the student and teacher network and three losses are evaluated. Firstly, we train the student network using click-level labels by minimizing the partial cross-entropy loss L_{pCE} [44], which is defined as

$$L_{pCE} = -\frac{1}{n} \sum_{i \in n} \hat{y}_i^c \log(y_i^c), \quad (2)$$

where $i \in n$ indicates that only labeled pixels participate in the calculation of the loss, and $\hat{y}_i^c = [0, 1]^c$ is the ground truth of pixel i belonging to class c .

To obtain the assistance of the ancillary teacher network, we apply a pixel consistency loss L_{pCons} , which is defined as

$$L_{pCons} = -\frac{1}{N} \sum_{i \in N} \|f(x_i, \theta') - f(x_i, \theta)\|^2, \quad (3)$$

where $f(\cdot)$ is the softmax prediction of the network and no gradient is calculated in the teacher network.

An regularized loss L_{CRF} [44] is also applied to smooth the segmentation, which is defined as

$$L_{CRF} = \sum_C Y^{C'} W_{pq} (1 - Y^C), \quad (4)$$

where W_{pq} is a dense Gaussian kernel with a role of the relaxation of dense CRF [19], Y^C is the softmax output of each class, and $Y^{C'}$ is the transposed matrix of Y^C .

After the backpropagation of all the losses, the ancillary teacher network will be updated by EMA. This process continues iteratively until the end of the training. Overall, the ancillary student network is trained by loss L^* , which is defined as

$$L^* = L_{pCE} + \lambda_{pCons} L_{pCons} + \lambda_{CRF} L_{CRF}, \quad (5)$$

where λ controls the contribution of each loss term.

Training the primary model. In the primary model, we also use the teacher-student module for the training. In addition, we apply the student-student module to connect the ancillary student network and the primary student network with heterogeneous pseudo-labels.

After the ancillary model is fully trained, the networks of the primary student and teacher are initialized in the same manner as their ancillary counterparts. During each training iteration, we train the primary student network the same way as the ancillary student network through EMA. Moreover, we input training images to the ancillary model, and obtain prediction maps. By choosing the maximal class of the prediction maps, we generate heterogeneous pseudo-labels to introduce the contribution of the ancillary student network to the training of the primary student network. To include the information of heterogeneous pseudo-labels, a new loss L_{pseudo} is proposed. Since heterogeneous pseudo-labels are applied to each pixel, the loss is in the form of cross-entropy. L_{pseudo} is defined as

$$L_{pseudo}(\theta) = -\frac{1}{N} \sum_{i \in N} \tilde{y}_i^c \log(y_i^c), \quad (6)$$

where \tilde{y} denotes the heterogeneous pseudo-labels generated by ancillary model θ_{anc} .

Therefore, the primary student network is trained with the overview loss L in each iteration, which is defined as

$$L = L^* + \lambda_{pseudo} L_{pseudo}. \quad (7)$$

3.2. Collaboration of modules

Teacher-student. A large number of unlabeled pixels are not well utilized in click-level supervised semantic segmentation, which is also the case in semi-supervised learning (SSL). Thus, we regard the click-level supervision as a SSL task, where some image pixels are labeled while the others are unlabeled. The mean-teacher [45] is a effective SSL method that use a teacher-student module to leverage unlabeled images. Inspired by this, we adapt the teacher-student module to our model by operating on unlabeled pixels instead of unlabeled images.

In the teacher-student module, the teacher network is obtained by the EMA of the student network. The EMA network is proved to be more efficient than using the final network directly [35]. EMA can be considered a temporal ensemble process, which endows it with a strong generalization ability. Thus, teacher network can avoid the overfitting to click-level labels and further guides the student network to learn the full object. In addition, its ability to reduce the bias of the targets can achieve a smoother classification boundary [45]. Since the object boundary can be viewed as classification boundary [14], the EMA network can also predict a smoother and more accurate mask.




Figure 3. Visualization of the mechanism in seminar learning. We obtain the first three results in the tenth epoch of primary model training.

To make consistency constraint between teacher and student network, we propose a pixel consistency loss L_{pCons} , as a form of mean square error (MSE). Our pixel consistency loss only measures unlabeled pixels and is defined as:

$$L_{pCons} = \frac{1}{N - n} \left(\sum_{i \in n} \|f(x_i, \theta') - f(x_i, \theta)\|^2 - \sum_{i \in N} \|f(x_i, \theta') - f(x_i, \theta)\|^2 \right) \quad (8)$$

Because $n \ll N$, we ultimately use an approximate form of L_{pCons} , defined in Eq. 3

Student-student. The teacher-student module is used as an individual model, while the student-student module is used to connect two models. In the student-student module, we propose heterogeneous pseudo-labels as the bridge between the two models. The heterogeneous pseudo-labels are generated by the ancillary student network after the network is fully trained. Then, the labels are transferred to the primary student network.

Early attempts with pseudo-labels [22] used the network’s predictions to train the network itself. However, such an operation will produce confirmation bias [2]. In this case, the model will memorize the false pseudo-labels and it will be difficult to forget them during training. Therefore, we use a fully trained ancillary model to generate heterogeneous pseudo-labels. Such a discriminative model can produce reliable predictions that guide the training of the primary student network correctly.

Furthermore, the ancillary model should be trained under different conditions from the primary model, such as different random seed. As is mentioned before, models can generate different masks with great diversity. Learn from the prediction of the ancillary student network can compensate the deficiency of the primary student network and then smooth segmentation boundary.

3.3. Mechanism of Seminar Learning

We visualize the prediction of each network in the training of click-level supervised semantic segmentation, as

shown in Fig. 3, to illustrate why seminar learning works by leveraging teacher-student and student-student modules.

Comparing the segmentation of the ancillary student and primary student, we can see that the two networks have diverse predictions of the target person. Although the ancillary student fails to predict the right arm of the person in the green box, it has better robustness to noise in the red box and correctly predicts the legs of the person in the yellow box. The limbs of the person and background noise are uncertain regions since they are far from the click-level labels. By integrating the two networks in the student-student module, the ancillary student obtain a better segmentation performance in the leg and noisy regions.

As for the teacher-student module, we find that the prediction of the primary teacher covers a wider region of the person in the green and yellow boxes compared to the primary student, which confirms that the primary teacher has better generalization. Since the primary teacher is updated by the primary student, the learning range of the primary student will gradually grow during training, finally allowing it to recognize the whole person.

After the training is done, we obtain the final segmentation prediction as the output. We can see that almost every part of the person is accurately predicted, and the final result is close to the ground truth. This shows that our seminar learning can effectively integrate information from all the networks in our pipeline and overcome the limitations of click-level labels to provide smoother segmentation.

4. Experiments

4.1. Experimental Setup

Implementation details. We choose the network DeepLab-v3+ [10] with ResNet101 [15] as our backbone due to its flexible structure and excellent performance. In each model, we train the student network with a batch-size of 12 over 60 epochs. We follow the optimization strategy in DeepLab-V3+ for the baseline with full supervision, using stochastic gradient descent (SGD) [7] with a base learning rate of 0.007, polynomial schedule [26], momen-

Method	Foreground Annotation	Background Annotation	Specifics	mIOU (%)
What’s the Point [4]	manual	-	VGG16, size=[1×1]px	43.40
ScribbleSub [24]	synthetic	synthetic	Deeplab-v2-VGG16, size=[3×3]px	51.60
Regularized Loss [44]	synthetic	synthetic	Deeplab-v2-ResNet101, size=[3×3]px	57.00
Regularized Loss [44]	manual	synthetic	Deeplab-v3+-ResNet101, size=[1×1]px	55.63
Ours	manual	synthetic	Deeplab-v3+-ResNet101, size=[1×1]px	72.51

Table 1. Performance comparison on the Pascal VOC 2012 validation set. For all methods, we report performance under click-level supervision. We also describe the training specifics and the annotation source in the table.

tum of 0.9, and weight decay of $5e^{-4}$ for 90K steps. We set the hyperparameters of the loss weights as $\lambda_{CRF} = 1$, $\lambda_{pseudo} = 1$ and $\lambda_{pCons} = 200$. These hyperparameters remain the same across all models. In each model, since the teacher network does not perform well in the early epochs, we ramp up the weight λ_{pCons} from 0 to its final value during the first 40 epochs. The λ_{pCons} is updated in each epoch with a linear function. We set the EMA decay parameter $\alpha = 0.999$ for each teacher-student module. Our implementation is based on Pytorch [34]. All experiments are run on an Nvidia Titan RTX (24G) GPU.

Datasets and annotations. The majority of our experiments are on the PASCAL VOC 2012 segmentation dataset [13], which contains 10,582 training images, 1,449 validation images and 1,456 test images. We use the same setting as previous click-level supervision works. The mean intersections over union (mIOU) averaged over 21 classes are evaluated on the validation set. In previous works, there are two methods to obtain click-level labels for the Pascal VOC 2012 dataset. ‘What’s the point’ [4] first proposed the concept of click-level supervised semantic segmentation, and manually annotated each object with only one pixel in an image. Each click-level label in [24, 44] is one point randomly selected from the scribble label of each object. We choose manual annotation as our training labels because they can better reflect the real click-level supervised annotations. However, the non-instance class (background) is not annotated in this annotation set. Thus we use the background labels proposed in [32], which are synthetically generated from scribble labels.

Data augmentation. Following [44], we only use the default data augmentation in DeepLab-V3+ to handle image data. It includes random scaling crop, horizontal flipping and random Gaussian noise for the training set and fixed scaling crop for the validation set. We set the image crop size to 513×513 for both the training and validation sets.

4.2. Main Results

Table 1 compares our method against the SOTA weakly supervised approaches on the Pascal VOC 2012 validation set. In the table, the two ‘annotation’ columns represent the sources of the click-level labels. ‘Manual’ indicates

that click-level labels are generated by manual annotation, and ‘synthetic’ indicates that they are synthetically generated from scribbled labels. As mentioned in Sec. 4.1, our annotation set consists of manually annotated foreground labels and background labels generated from the scribble. The first three rows of the table are results published in the corresponding papers. In order to compare with the current SOTA method under the same conditions, we run the ‘Regularized Loss’ method [44] on the same annotation set, backbone and label size as our method. Its result is shown in the fourth row, which achieves 55.63% mIOU when using the hyperparameter settings in its paper.

Our method is based on the implementation of ‘Regularized Loss’, which includes L_{pCE} and L_{CRF} . Without modifying any parameters of the implementation in the fourth row, we apply our seminar learning method directly to ‘Regularized Loss’. After this operation, we greatly improve the mIOU from 55.63% to 72.51% (increased by 16.88%), which far exceeds all previous SOTA methods.

Supervision Level	Annotation Time (sec/img)	Method	mIOU (%)
Box	38.1	BCM [42]	70.2
Scribble	34.9	BPG [47]	73.2
Image	20.0	DRS [17]	71.2
Click	22.1	Ours	72.5

Table 2. Performance comparison on the Pascal VOC 2012 validation dataset with various SOTA weakly supervised methods.

Comparison with other weakly supervision. In Table 2, we compare various levels of weakly supervised methods, all of which use ResNet101 as the backbone and no post-processing. The annotation times in the table are provided in [4, 5]. It can be seen that click-level supervision only acquires 22.1 seconds to annotate an image, which is close to image-level supervision. Meanwhile, our method gain a medial performance, which exceeds the image-level but is lower than the scribble-level, indicating that we make full use of the potential of click-level supervision and our method is a trade-off between time cost and performance.

L_{pCE}	L_{CRF}	L_{pCons}	L_{pseudo}	mIOU (%)	
				MobileNet	ResNet101
✓				50.85	54.70
✓	✓			48.75	55.63
✓		✓		56.14	61.47
✓	✓	✓		62.03	70.29
✓	✓	✓	✓	64.44	72.51
Full supervision				71.92	78.59

Table 3. Performance comparison of our framework on PASCAL VOC 2012 val set. The experiments aim to verify the effects of each module in seminar learning on different networks. For all modules we report performance both under click-level supervision

4.3. Ablation Study

In this section, we verify seminar learning by testing the effectiveness of losses, student-student module, and hyper-parameters separately.

The effect of the losses in seminar learning. This part aims to show the contributions of different losses on seminar learning. Table 3 shows the mIOU scores of five combinations of losses for comparison. Notably, teacher network has the similar performance with student network after training, and we use student network as the output. In the combination without L_{pseudo} , the output of the ancillary student is final prediction.

In fact, ‘regularized loss’ [44] is the combination of $L_{pCE} + L_{CRF}$. By comparing 1st and 2nd rows of Table 3, $L_{CRF} + L_{pCE}$ cannot improve the segmentation performance evidently under the click-level supervision.

By comparing 2nd-4th rows of Table 3, we can find that $L_{pCE} + L_{CRF} + L_{pCons}$ provides a significant improvement where the mIOU grows from 55.63% to 70.29%, increasing by 14.66%. Its result also exceeds that of $L_{pCE} + L_{pCons}$ by 8.82%, which indicates that L_{CRF} is helpful for improving the performance by cooperating with L_{pCons} . The 5th row shows that adding L_{pseudo} to the training can receive an extra 2.22% improvement and reach up to 72.51%.

Therefore, all the losses in seminar learning contribute to the improvement of the mIOU in click-level weakly supervised segmentation. Although adding L_{CRF} to L_{pCE} is not effect, L_{CRF} can play a complementary role where intertwined with the pixel consistency loss L_{pCons} .

To verify the generality of the proposed losses, we also test the losses on Deeplab-V3+ with MobileNet. The results of MobileNet have the same trend as that of ResNet101. It can prove that our losses can be used in other networks.

The effectiveness of the student-student module. In student-student module, heterogeneous pseudo-labels are specially designed to transfer between two different networks. As shown in Fig. 4, we compare our module with two other cases to verify the effectiveness of student-student




Figure 4. Comparison of different methods that use pseudo-labels on the Pascal VOC 2012 validation set. All results are obtained by DeeplabV3+ with ResNet101.

module. In the compared cases, the pseudo-labels are transferred to the same student network instead of a different student network, which we called self-generation method. Two cases use different learning rate schemes.

In the first 60 epochs, the three cases use the same ancillary model. We generate pseudo-labels in the 61th epoch. After the 60th epoch, we apply two learning rate schemes for the self-generation method: 1) Self-Generation-reset: we reset the learning rate at 61th epoch and keep the learning rate consistent with the student-student. It is obvious that the mIOU of the this case is always lower than that of the student-student module after training with pseudo-labels. The result becomes worse than that before the pseudo-labels are used. 2) Self-Generation-unchange: the learning rate remains unchanged after the 60th epoch. mIOU remains stable until the end of training. From the experiment, we find that the pseudo-labels generated by self-generation cannot improve the model. The student-student module is proved to be necessary to improve click-level supervised segmentation.

The number of student-student modules. The fully trained primary model can also act as the ancillary model for the next primary model. Therefore, our method can extend to a chain of student-student modules. Fig. 5(a) shows the performance in different numbers of student-student modules. We use two networks to evaluate how many student-student modules can improve the performance. When applying the first student-student module, the performance of the network is improved significantly. The accuracy of the Resnet101 improves from 70.29% to 72.51%, while that of the MobileNet improves from 60.59% to 62.59%. When we apply student-student modules more than once, the mIOU in both two networks almost does not increase. It concludes that one student-student module is




Figure 5. Ablation studies and empirical analysis on the Pascal VOC 2012 validation set. (a) Validation mIOU (%) with different numbers of student-student modules; (b) Validation mIOU over the EMA decay weight α ; (c) Validation mIOU over weight λ_{pCons} of the pixel consistency loss.




Figure 6. Visualization on the Pascal VOC 2012 validation set.

sufficient for the training of networks.

Sensitivity analysis of hyperparameters. The performance of seminar learning depends on the hyperparameters of consistency weight λ_{pCons} and EMA decay α .

Fig. 5(b) shows the sensitivity of decay weight α in the EMA. From [45] we know that good decay always spans roughly an order of magnitude, and the commonly used hyperparameter values in current semi-supervised methods are 0.999 and 0.99. Our experiments prove that 0.999 is better in the click-level supervision task.

Fig. 5(c) shows the sensitivity of the network to the consistency weight λ_{pCons} . We find that the value of λ_{pCons} is much smaller than λ_{pCE} of click-level supervision. This does not affect the training of the neural network. Thus, we increase the weight of the consistency loss L_{pCons} to reach the same order of magnitude as L_{pCE} . Using this order of magnitude, we experiment with different hyperparameters on the two networks. We find that the student network can learn the best results from the teacher network when the consistency weight reaches 200.

4.4. Visualization

Fig. 6 shows examples of semantic segmentation predictions from qualitative views. In the first four columns, we apply different combinations of losses with seminar learning, and the segmentation results show the effects of them.

Visualized explanation of losses. As indicated in Sec. 4.3, L_{CRF} contributes to the performance only when it is combined with the pixel consistency loss L_{pCons} . Based on the theory of dense CRF [18], if two pixels are close in both color and distance, the category association between them will be tight, and the two pixels are more likely to be predicted as the same category. Thus, the dense CRF based loss L_{CRF} can make the segmentation result continuous inside objects. From the first two columns of Fig. 6, we can observe that the addition of L_{CRF} enhances the internal continuity of the same category. However, as the limited labels contains insufficient information, the model fails to predict the pixels near the boundary of the object. Since L_{CRF} is calculated with the prediction maps of the model, a large number of wrong predictions in the object prevent L_{CRF} from establishing the correct connection among pixels, so L_{CRF} does not work. In the third column of Fig. 6, after L_{pCons} is applied, the boundary prediction is more accurate, and L_{CRF} can smooth the predictions in the right way. Ultimately, our seminar learning achieve excellent performance with the combination of all losses.

5. Conclusion

In this paper, we propose a seminar learning paradigm for click-level weakly supervised semantic segmentation. Our approach consists of teacher-student and student-student modules, where we aggregate the generalized and diverse information from multiple networks. In this way, we address the insufficient information of the limited click-level labels. The experimental results demonstrate that our method achieves an excellent segmentation performance, surpassing the current SOTA model by 16.88% (mIOU).

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4981–4990, 2018. 3
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. 5
- [3] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017. 2
- [4] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. 2016. 1, 2, 3, 6
- [5] Míriam Bellver Bueno, Amaia Salvador Aguilera, Jordi Torres Viñals, and Xavier Giró Nieto. Budget-aware semi-supervised semantic and instance segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, 2019*, pages 93–102, 2019. 6
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5049–5059, 2019. 2
- [7] Léon Bottou. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012. 5
- [8] Liyi Chen, Weiwei Wu, Chenchen Fu, Xiao Han, and Yuntao Zhang. Weakly supervised semantic segmentation with boundary exploration. In *European Conference on Computer Vision*, pages 347–362. Springer, 2020. 3
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1, 2
- [10] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 2, 5
- [11] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 6830–6840, 2019. 2
- [12] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1635–1643, 2015. 1, 3
- [13] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 6
- [14] Geoffrey French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations. In *British Machine Vision Conference*, number 31, 2020. 4
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [16] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017. 3
- [17] Beomyoung Kim, Sangeun Han Kim, et al. Discriminative region suppression for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2103.07246*, 2021. 3, 6
- [18] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, pages 109–117, 2011. 8
- [19] Philipp Krähenbühl and Vladlen Koltun. Parameter learning and convergent inference for dense random fields. In *International Conference on Machine Learning*, pages 513–521. PMLR, 2013. 4
- [20] Viveka Kulharia, Siddhartha Chandra, Amit Agrawal, Philip Torr, and Amrith Tyagi. Box2seg: Attention weighted loss and discriminative feature learning for weakly supervised segmentation. In *European Conference on Computer Vision*, pages 290–308. Springer, 2020. 3
- [21] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. 2
- [22] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 2013. 2, 5
- [23] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5267–5276, 2019. 3
- [24] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 3, 6
- [25] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017. 2
- [26] Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. *arXiv preprint arXiv:1506.04579*, 2015. 5

- [27] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [28] Wenfeng Luo and Meng Yang. Semi-supervised semantic segmentation via strong-weak dual-branch network. In *European Conference on Computer Vision*, pages 784–800. Springer, 2020. 3
- [29] Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, and Bo Zhang. Smooth neighbors on teacher graphs for semi-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8896–8905, 2018. 2
- [30] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *arXiv preprint arXiv:2001.05566*, 2020. 1
- [31] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 2
- [32] Anton Obukhov, Stamatios Georgoulis, Dengxin Dai, and Luc Van Gool. Gated crf loss for weakly supervised semantic image segmentation. *arXiv preprint arXiv:1906.04651*, 2019. 6
- [33] George Papandreou, Liang Chieh Chen, Kevin Murphy, and Alan L. Yuille. Weakly- and semi-supervised learning of a dcnn for semantic image segmentation. 2015. 1
- [34] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [35] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992. 4
- [36] Rui Qian, Yunchao Wei, Honghui Shi, Jiachen Li, Jiaying Liu, and Thomas Huang. Weakly supervised scene parsing with point-based distance metric learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8843–8850, 2019. 3
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2
- [38] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ” grabcut” interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 23(3):309–314, 2004. 3
- [39] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 2
- [40] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020. 2
- [41] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3136–3145, 2019. 3
- [42] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 6
- [43] Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, and Christopher Schroers. Normalized cut loss for weakly-supervised cnn segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1818–1827, 2018. 3
- [44] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 507–522, 2018. 2, 3, 4, 6, 7
- [45] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. 2017. 2, 4, 8
- [46] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. *arXiv preprint arXiv:1903.03825*, 2019. 2
- [47] Bin Wang, Guojun Qi, Sheng Tang, Tianzhu Zhang, Yunchao Wei, Linghui Li, and Yongdong Zhang. Boundary perception guidance: A scribble-supervised semantic segmentation approach. In *IJCAI*, pages 3663–3669, 2019. 6
- [48] Man Zhang, Yong Zhou, Jiaqi Zhao, Yiyun Man, and Rui Yao. A survey of semi- and weakly supervised semantic segmentation of images. *Artificial Intelligence Review*, 53(2), 2019. 1
- [49] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 3