

Adversarial Learning for Semi-Supervised Semantic Segmentation

Supplementary Materials

1 Pixel Accuracy in Semi-Supervised Learning

In Table 1, we show the average segmentation accuracy with respect to the number of selected pixels based on different threshold values of T_{semi} as in (5) of the paper on the Cityscapes dataset. With a higher T_{semi} , the discriminator outputs are more confident (similar to ground truth label distributions) and lead to more accurate pixel predictions. Also, as a trade-off, the higher threshold (T_{semi}), the fewer pixels are selected for back-propagation. This trade-off could also be observed in Table 5 of the paper.

Table 1: Selected pixel accuracy.

T_{semi}	Selected Pixels (%)	Accuracy
0	100%	92.65%
0.1	36%	99.84%
0.2	31%	99.91%
0.3	27%	99.94%

2 Additional Hyper-parameter Analysis

In Table 2, we show the complete hyper-parameter analysis. In addition to the analysis of λ_{semi} and T_{semi} in Table 5 of the paper, we show that the proposed adversarial learning is also robust to different values of λ_{adv} .

3 Training Parameters

In Table 3, we show the training parameters for both datasets. We use the PyTorch implementation, and we will release our code and models for the public.

Table 2: Hyper parameter analysis.

Data Amount	λ_{adv}	λ_{semi}	T_{semi}	Mean IU
Full	0	0	N/A	73.6
Full	0.005	0	N/A	74.0
Full	0.01	0	N/A	74.9
Full	0.02	0	N/A	74.6
Full	0.04	0	N/A	74.1
Full	0.05	0	N/A	73.0
1/8	0.01	0	N/A	67.6
1/8	0.01	0.05	0.2	68.4
1/8	0.01	0.1	0.2	69.5
1/8	0.01	0.2	0.2	69.1
1/8	0.01	0.1	0	67.2
1/8	0.01	0.1	0.1	68.8
1/8	0.01	0.1	0.2	69.5
1/8	0.01	0.1	0.3	69.2
1/8	0.01	0.1	1.0	67.6

Table 3: Training parameters.

Parameter	Cityscaps	PASCAL VOC
Trained iterations	40,000	20,000
Learning rate	2.5e-4	2.5e-4
Learning rate (D)	1e-4	1e-4
Polynomial decay	0.9	0.9
Momentum	0.9	0.9
Optimizer	SGD	SGD
Optimizer (D)	Adam	Adam
Nesterov	True	True
Batch size	2	10
Weight decay	0.0001	0.0001
Crop size	512x1024	321x321
Random scale	No	Yes

4 Additional Qualitative Results

In Figure 1-2, we show additional qualitative comparisons with the models using half training data of the PASCAL VOC dataset. In Figure 3, we show more qualitative comparisons with the models using half training data of the Cityscapes dataset. The results show that both the adversarial learning and the semi-supervised training scheme improve the segmentation quality.




Figure 1: Comparisons on the PASCAL VOC dataset using 1/2 training data.



Figure 2: Comparisons on the PASCAL VOC dataset using 1/2 training data.



Figure 3: Comparisons on the Cityscapes dataset using 1/2 training data.