

SEMI-SUPERVISED SEMANTIC SEGMENTATION VIA BOOSTING UNCERTAINTY ON UNLABELED DATA

Anonymous authors

Paper under double-blind review

1 SUPPLEMENTARY MATERIAL

1.1 MORE IMPLEMENTATION DETAILS

We implement our method based on the PyTorch framework, and all experiments are done with 4 NVIDIA V100 GPUs. The backbones of the two networks are initialized with the same ImageNet pre-trained weights, while the two Deeplabv3+ modules are randomly initialized. Following CPS (Chen et al., 2021), the tradeoff weight λ is set to 1.5 for PASCAL VOC 2012 and 6.0 for Cityscapes. We only use simple and traditional data augmentation: mirroring, resize, and crop. We employ a poly learning rate policy where the initial learning rate is multiplied by $(1 - \frac{iter}{max_iter})^{0.9}$.

1.2 QUALITATIVE RESULTS

We visualize the prediction maps of the baseline method CPS (Chen et al., 2021) and our method on PASCAL VOC 2012 and Cityscapes in Figures 1, 2, and 3.

1.2.1 RESULTS OF PASCAL VOC 2012

In Figure 1 and 2, we can see that, our result outperform baseline in various conditions. We also notice that our method can effectively reduce the uncertainty for those areas with high uncertainty in baseline; for those areas with wrong segmentation, our method can increase the uncertainty to better segment the image. These phenomenons prove that the model can achieve better generalization due to our strategies of boosting uncertainty on unlabeled images.

1.2.2 RESULTS OF CITYSCAPES DATASETS

In figure 3, we can see that our method produces a more consistent and smooth segmentation map compared with the baseline. Due to the region-aware nature of our uncertainty booster, our method can progressively enlarge the correct region and gradually cover a large portion of the target region, pushing the final mIoU higher. We specially highlight the details by yellow boxes in Figure 3. Similar to the results in PASCAL VOC 2012, our results are much cleaner, and our booster suppresses some noticeable mistakes, e.g., the wood in line 4 and the wall in line 6.

REFERENCES

Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2613–2622, 2021.

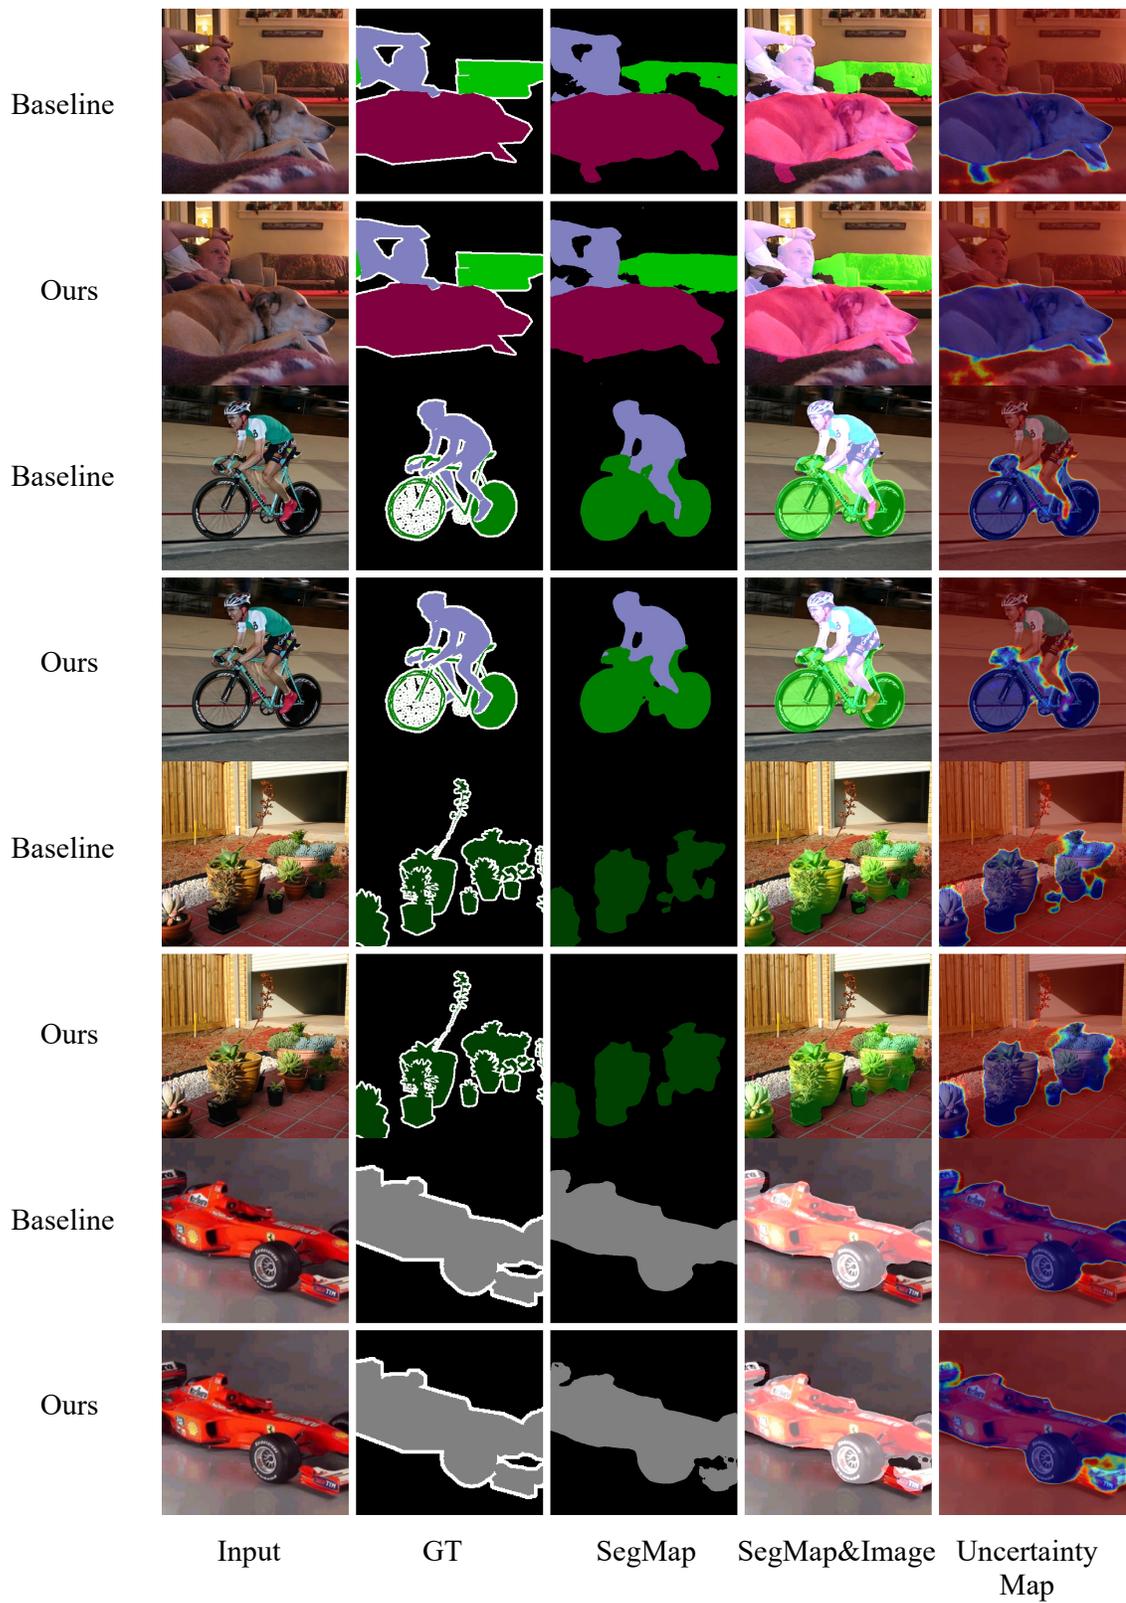


Figure 1: Visualization results on the PASCAL VOC 2012 val set. From top to bottom: baseline, our model ;From left to right: input, Ground Truth, segmentation map, segmentation map with input, uncertainty map.

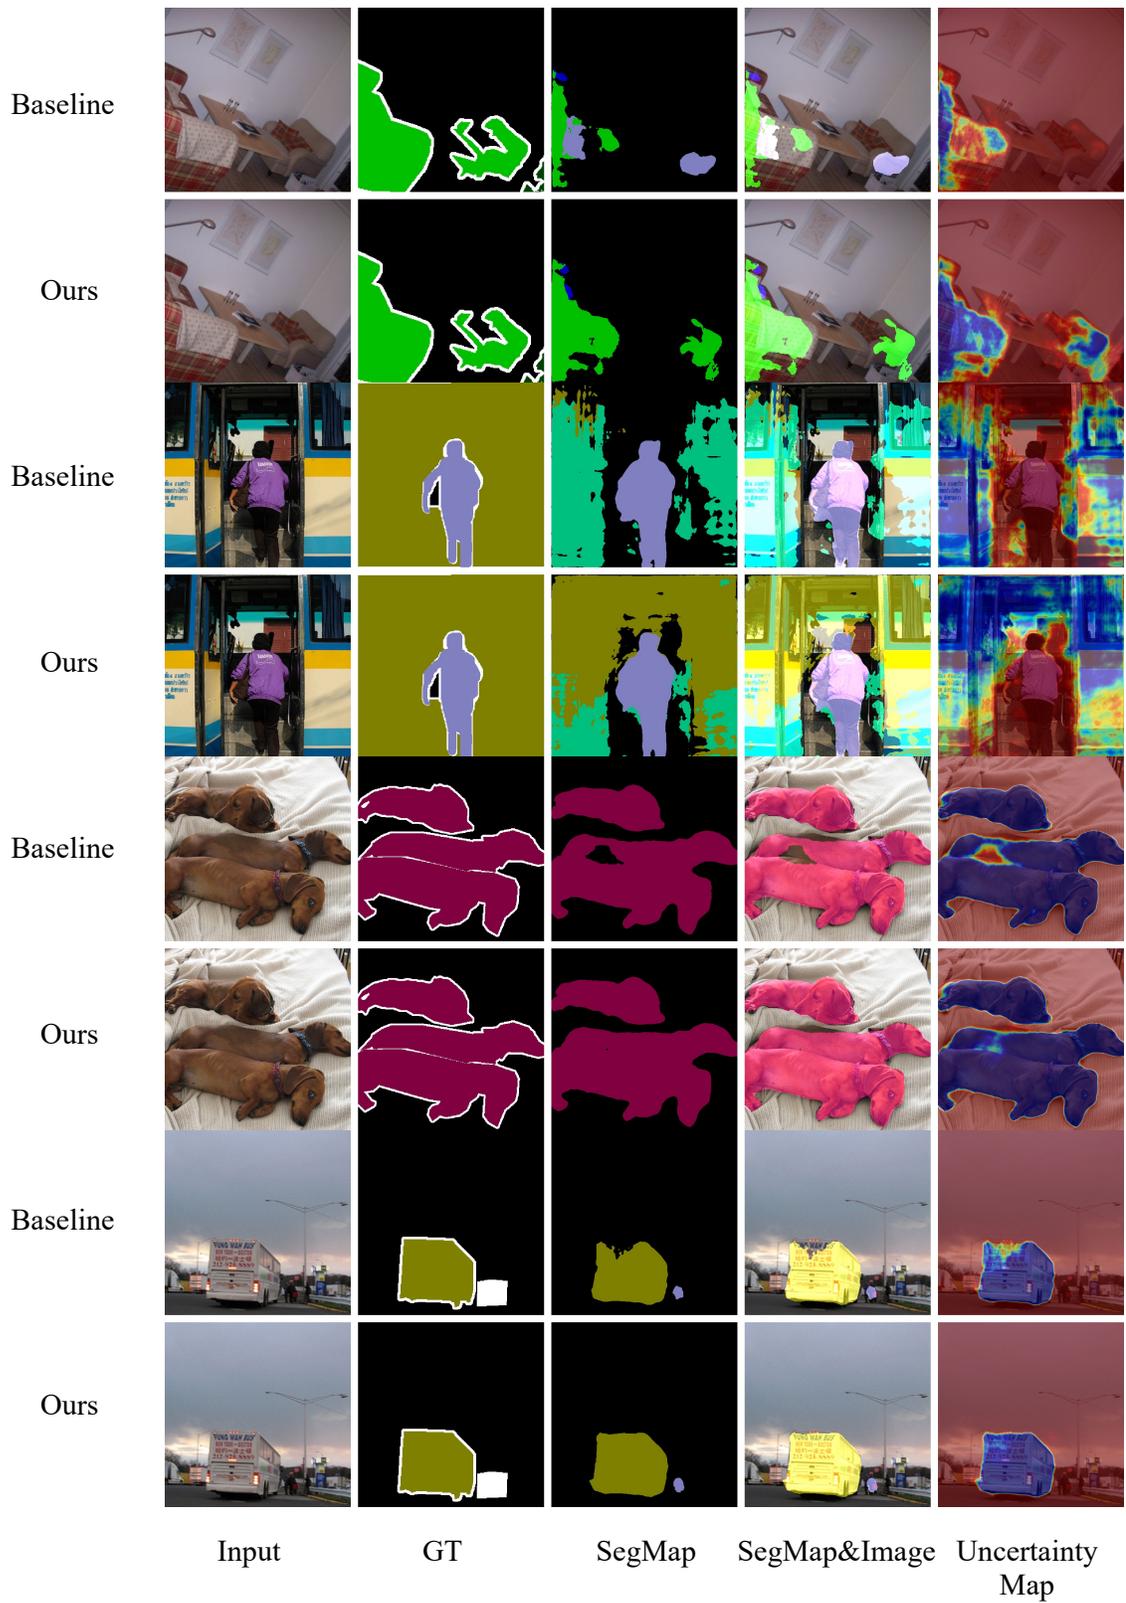


Figure 2: Visualization results on the PASCAL VOC 2012 val set. From top to bottom: baseline, our model ;From left to right: input, Ground Truth, segmentation map, segmentation map with input, uncertainty map.

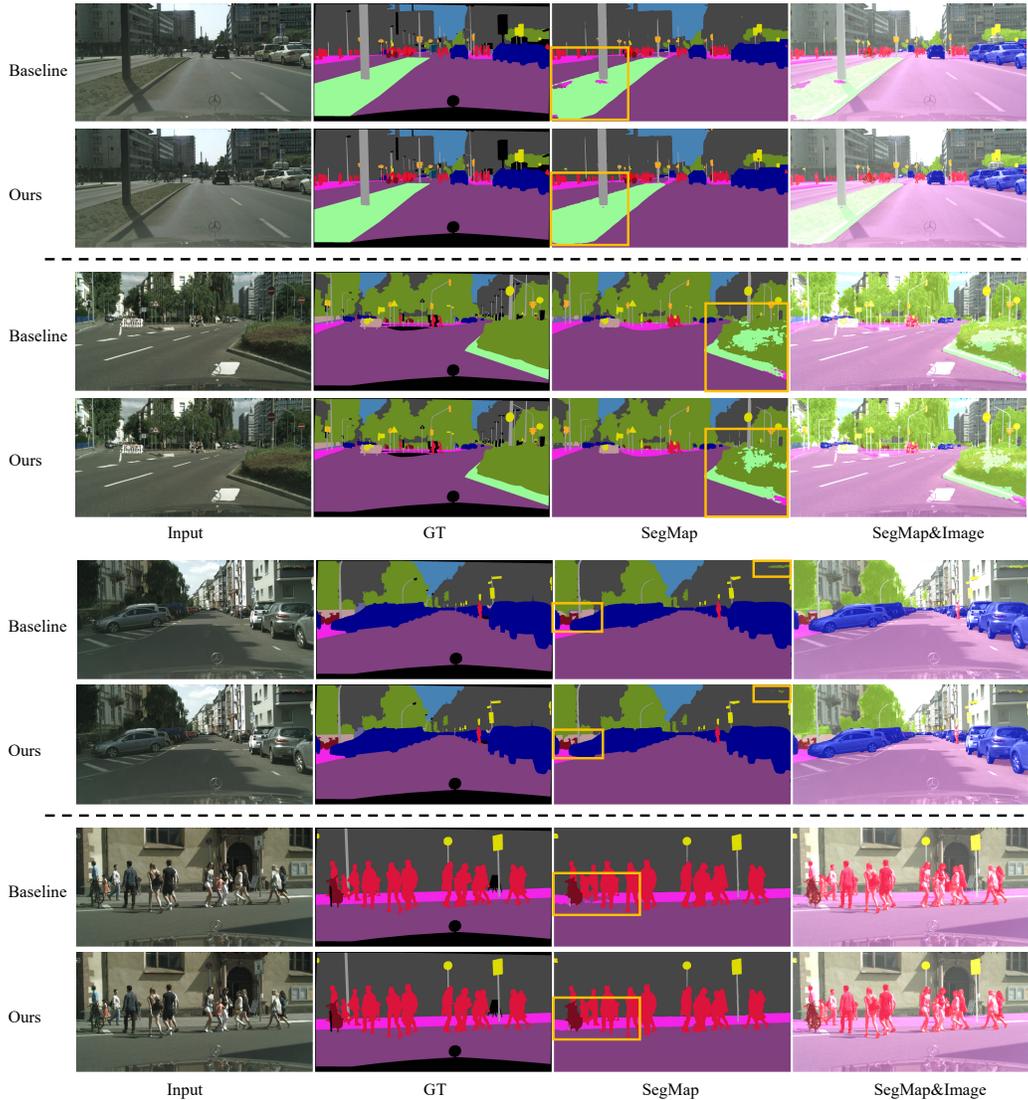


Figure 3: Visualization results on the Cityscapes val set. Zoom in for a clearer view. From top to bottom: baseline, our model; From left to right: input, Ground Truth, segmentation map, segmentation map with input image.