

Semantic Prompt Learning for Weakly-Supervised Semantic Segmentation

Supplementary Material

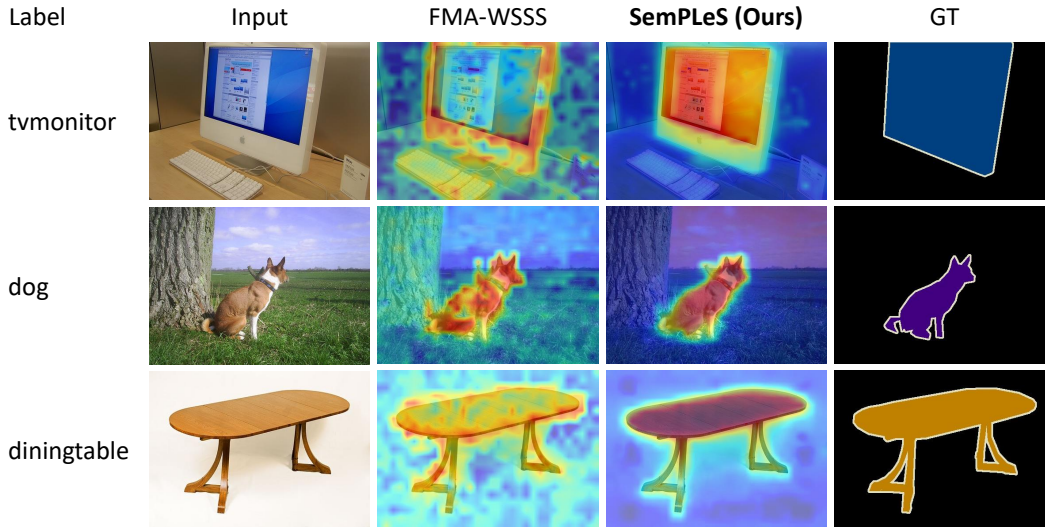


Figure 7. Qualitative results of CAMs. “GT” denotes the ground truth masks.

K	1	5	10	30	50
mIoU	68.2	68.4	68.5	68.7	68.4

Table 6. Quantitative analysis of the number of prompts.

λ	0	0.01	0.02	0.05	0.1
mIoU	67.6	68.3	68.4	68.7	67.8

Table 7. Quantitative analysis of loss weight.

Qualitative results: In Figure 7, we see that our proposed SemPLES framework results in more accurate CAMs which properly capture the whole target object and are well-aligned with the ground truth masks, demonstrating the effectiveness of our method.

6. More Experiments

The number of prompts: In Table 6, we provide analysis on the number of prompts K , and we see that the mIoU would stop increasing at $K = 50$ due to possible overfitting. Hence, we set $K = 30$ by default.

Loss weight: In Table 7, we provide analysis on the refinement loss weight λ . When varying λ , the mIoU gradually increases from 67.6% to 68.7% when $\lambda = 0.05$, and then drops to 67.8% when $\lambda = 0.1$. Hence, we set $\lambda = 0.05$ by default.