Progressive Mirror Detection Supplementary Material

Anonymous CVPR submission

Paper ID 9689

In this supplemental, we first present more examples from our proposed benchmark in Figure 1. We then show more qualitative results on both MSD and our benchmark test set. Specially, we compare our results with those of shadow detection methods DSC [4] and BDRAR [13]; semantic segmentation method PSPNet [11]; salient object detection methods R^3Net [3], CPDNet [9] and BASNet [7]; and mirror detection method MirrorNet [10] on the MSD dataset [10] and on our benchmark dataset (which is composed of images from ADE20K [12], COCO-Stuff [2], NYUD-V2 [6], Pascal-Context [5], SUNRGBD [8] and MINC [1]).

Noted that for the MSD dataset, we train our model and other baselines on its training dataset and test on its test dataset (Figure 2). For our benchmark dataset, we train on the other 5 datasets in our benchmark dataset, excluding the one that we want to evaluate, and then we test on the excluded one (Figure 3, 4, 5, 6, 7, 8).

References

[1] Sean Bell, Paul Upchurch, Noah Snavely, and Kavita Bala. Material recognition in the wild with the materials in context database. *CVPR*, 2015. 1, 4, 5, 6, 7

[2] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, June 2018. 1, 4, 5, 6, 7

[3] Zijun Deng, Xiaowei Hu, Lei Zhu, Xuemiao Xu, Jing Qin, Guoqiang Han, and Pheng-Ann Heng. R3net: Recurrent residual refinement network for saliency detection. In *IJCAI*, pages 684–690. AAAI Press, 2018. 1, 3, 4, 5, 6, 7

[4] Xiaowei Hu, Lei Zhu, Chi-Wing Fu, Jing Qin, and Pheng-Ann Heng. Direction-aware spatial context features for shadow detection. In *CVPR*, pages 7454–7462, 2018. 1, 3, 4, 5, 6, 7

[5] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, 2014. 1, 4, 5, 6, 7

[6] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 1, 4, 5, 6, 7

[7] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *CVPR*, June 2019. 1, 3, 4, 5, 6, 7

[8] Shuran Song, Samuel P. Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, June 2015. 1, 4, 5, 6, 7

[9] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *CVPR*, June 2019. 1, 3, 4, 5, 6, 7

[10] Xin Yang, Haiyang Mei, Ke Xu, Xiaopeng Wei, Baocai Yin, and Rynson W.H. Lau. Where is my mirror? In *ICCV*, October 2019. 1, 3, 4, 5, 6, 7

[11] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 1, 3, 4, 5, 6, 7

[12] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. 1, 4, 5, 6, 7

[13] Lei Zhu, Zijun Deng, Xiaowei Hu, Chi-Wing Fu, Xuemiao Xu, Jing Qin, and Pheng-Ann Heng. Bidirectional feature pyramid network with recurrent attention residual modules for shadow detection. In *ECCV*, 2018. 1, 3, 4, 5, 6, 7

Figure 1: More examples in our proposed benchmark dataset.

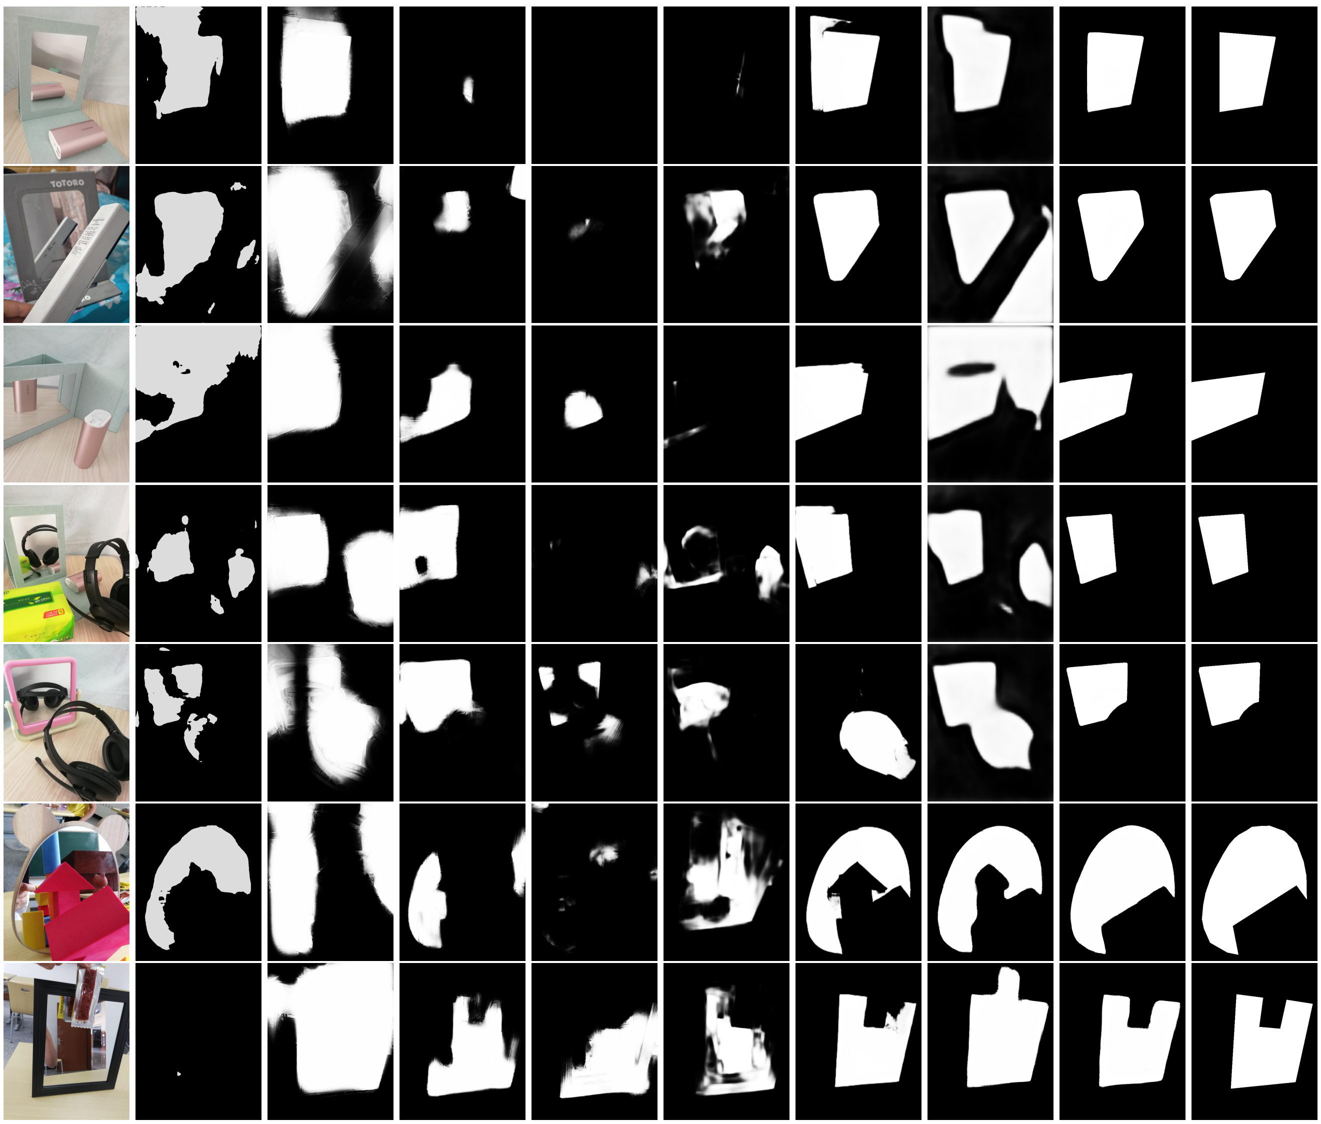

Figure 2: Qualitative results of our method, compared with relevant state-of-the-art methods on the MSD [10] test set. All methods are trained on the MSD training set.

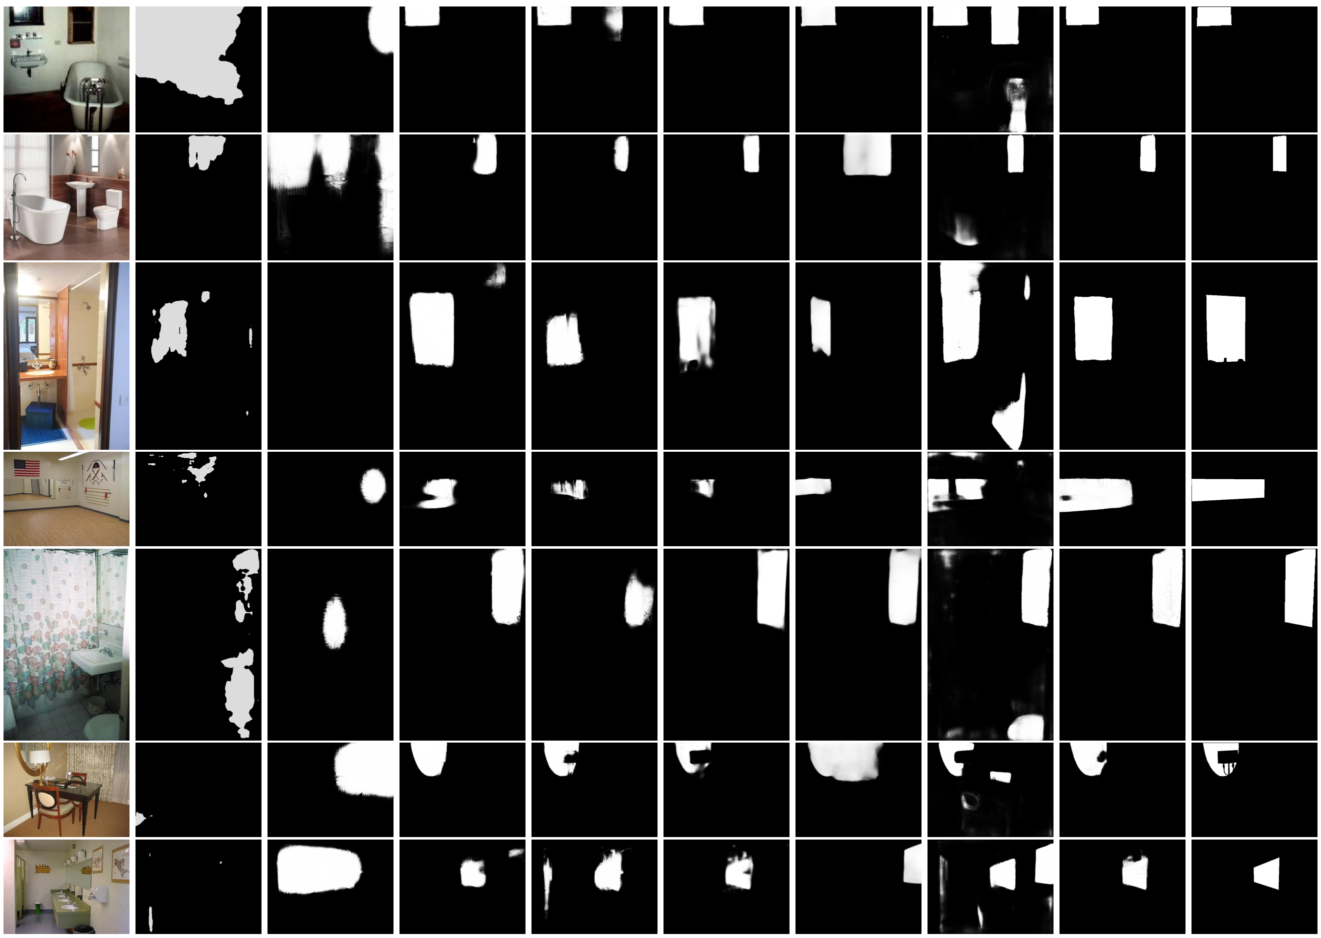

Figure 3: Qualitative results of our method, compared with relevant state-of-the-art methods on the ADE20K [12] dataset. All methods are trained on the COCO-Stuff [2], MINC [1], NYUD-V2 [6], Pascal-Context [5], and SUNRGBD [8] datasets.

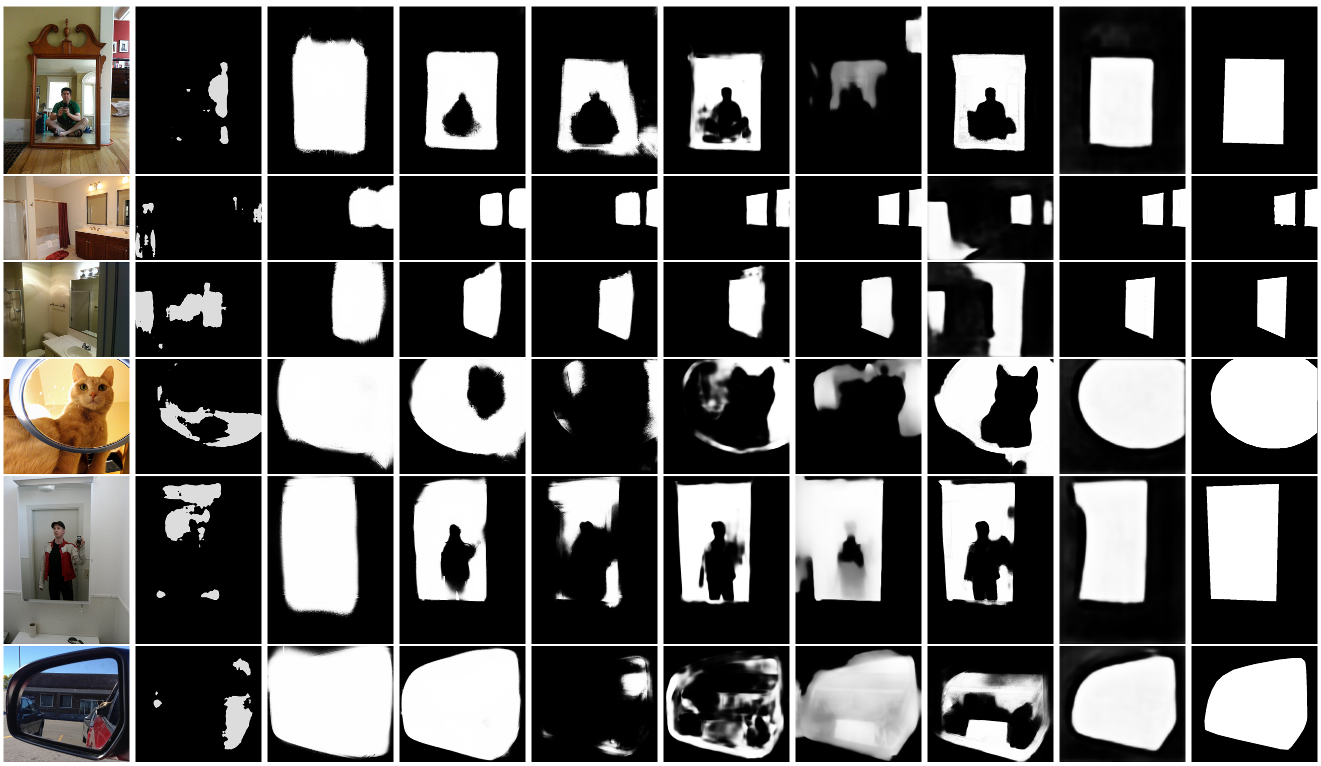

Figure 4: Qualitative results of our method, compared with relevant state-of-the-art methods on the COCO-Stuff [2] dataset. All methods are trained on the ADE20K [12], MINC [1], NYUD-V2 [6], Pascal-Context [5], and SUNRGBD [8] datasets.

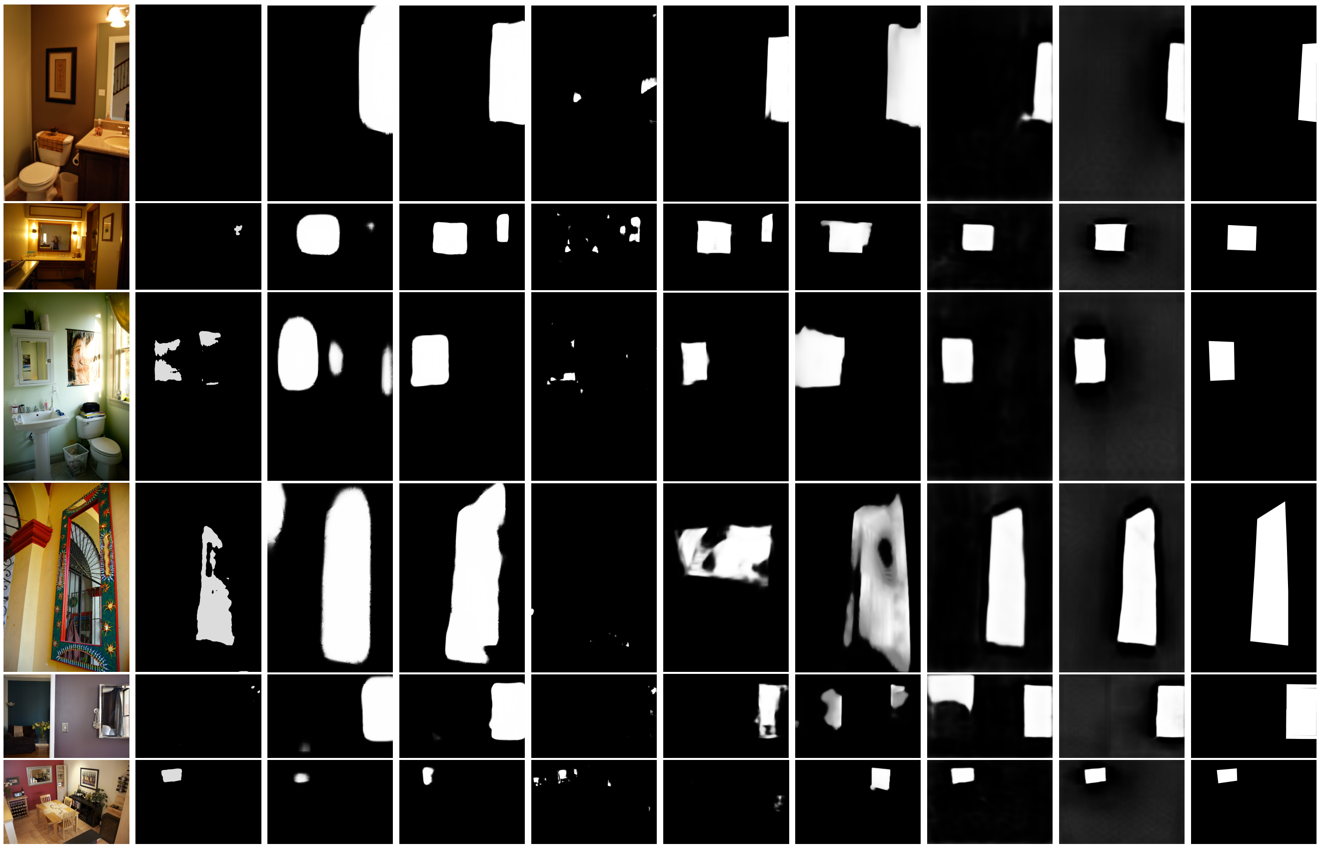

Figure 5: Qualitative results of our method, compared with relevant state-of-the-art methods on the MINC [1] dataset. All methods are trained on the ADE20K [12], COCO-Stuff [2], NYUD-V2 [6], Pascal-Context [5], and SUNRGBD [8] datasets.

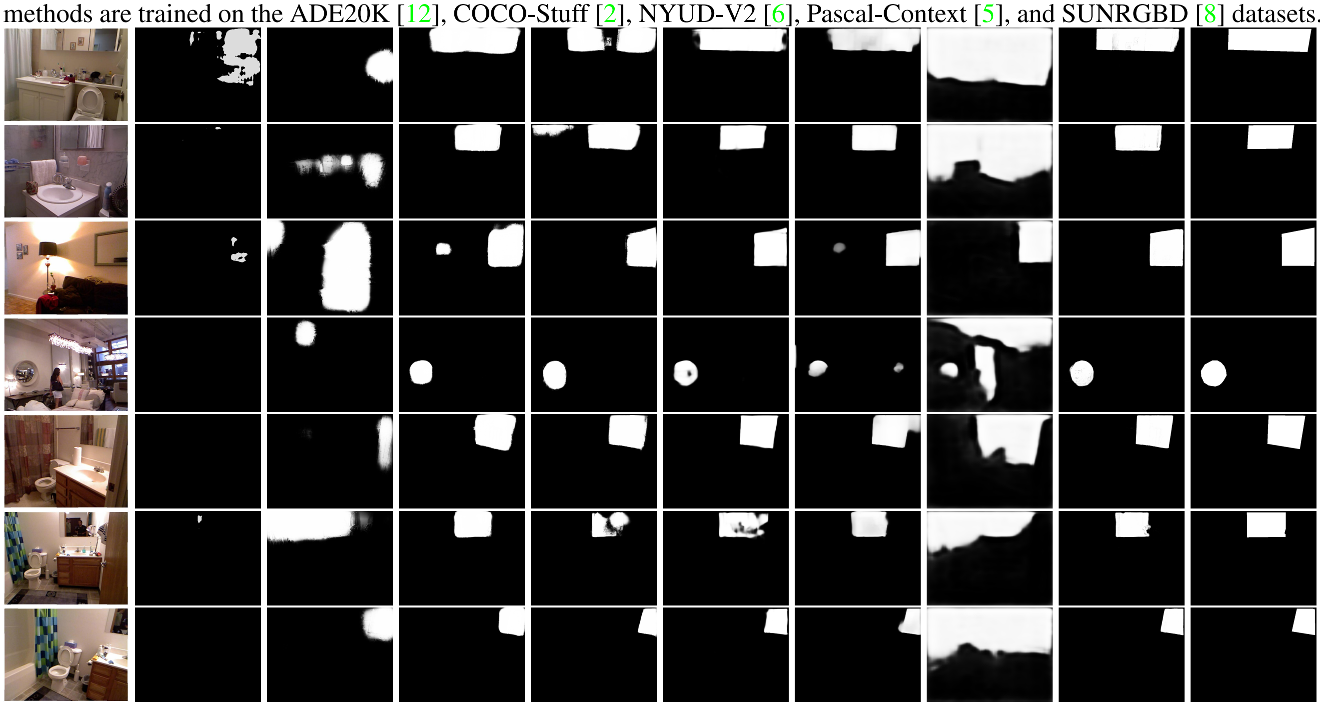

Figure 6: Qualitative results of our method, compared with relevant state-of-the-art methods on the NYUD-V2 [6] dataset. All methods are trained on the ADE20K [12], COCO-Stuff [2], MINC [1], Pascal-Context [5], and SUNRGBD [8] datasets.

CVPR #9689

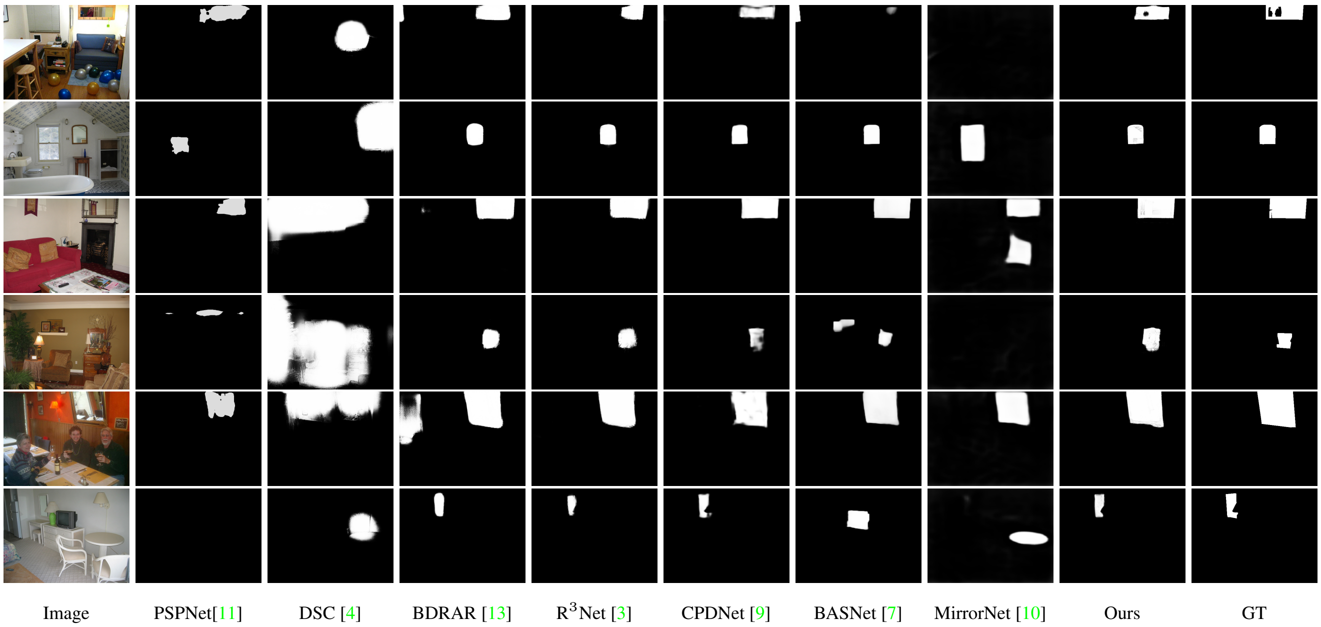

Figure 7: Qualitative results of our method, compared with relevant state-of-the-art methods on the Pascal-Context [5] dataset. All methods are trained on the ADE20K [12], COCO-Stuff [2], MINC [1], NYUD-V2 [6], and SUNRGBD [8] datasets.

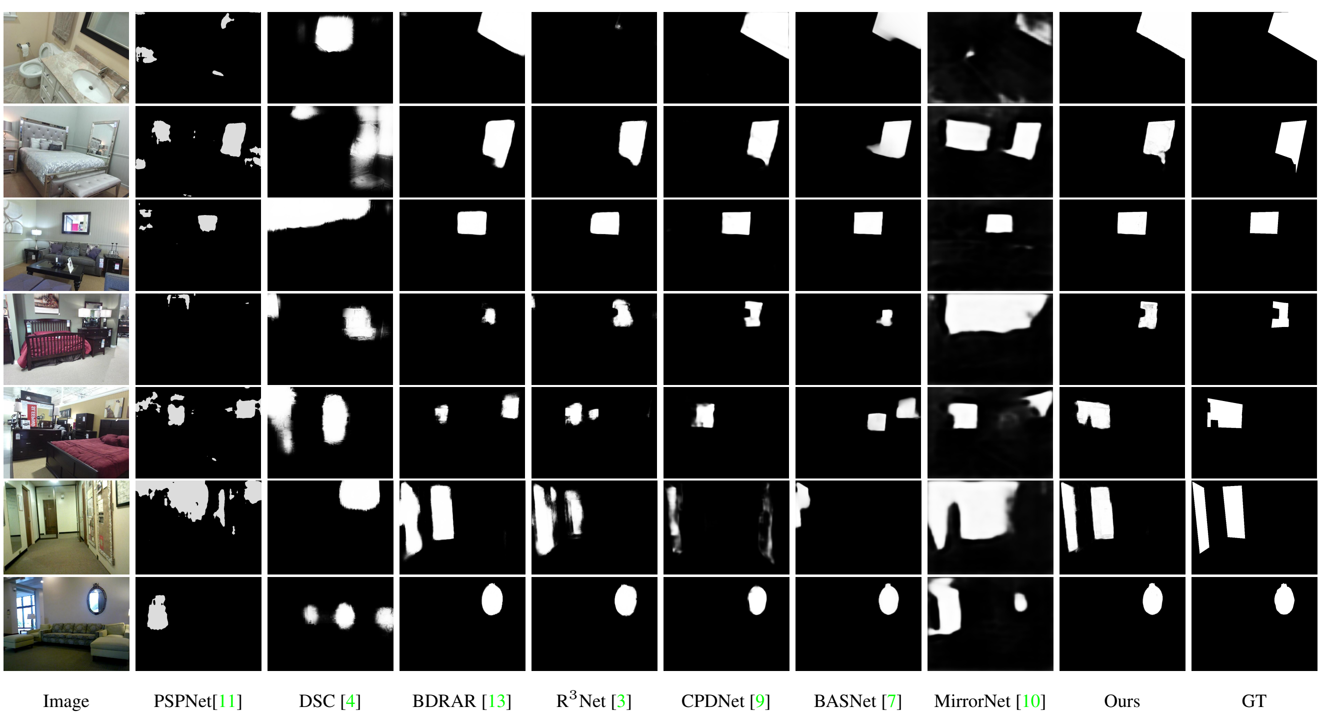

Figure 8: Qualitative results of our method, compared with relevant state-of-the-art methods on the SUNRGBD [8] dataset. All methods are trained on the ADE20K [12], COCO-Stuff [2], MINC [1], NYUD-V2 [6], and Pascal-Context [5] datasets.