Distilling Reflection Dynamics for Single-Image Reflection Removal
Supplementary Material

Quanlong Zheng[1] Xiaotian Qiao[2] Ying Cao[2] Shi Guo[3,4] Lei Zhang[3,4] Rynson Lau[2]
OPPO Research [2]City Unversity of Hong Kong
[3]The Hong Kong Polytechnic University [4]DAMO Academy, Alibaba Group

1. Introduction

We first explore the effect of input image sequence length in Section 2 and explore the impact of the synthetic and real datasets in Section 3. We then show some samples of our synthetic dataset and our newly collected real-world dataset in Section 4. Finally, we show more visual comparisons of different methods in Section 5.

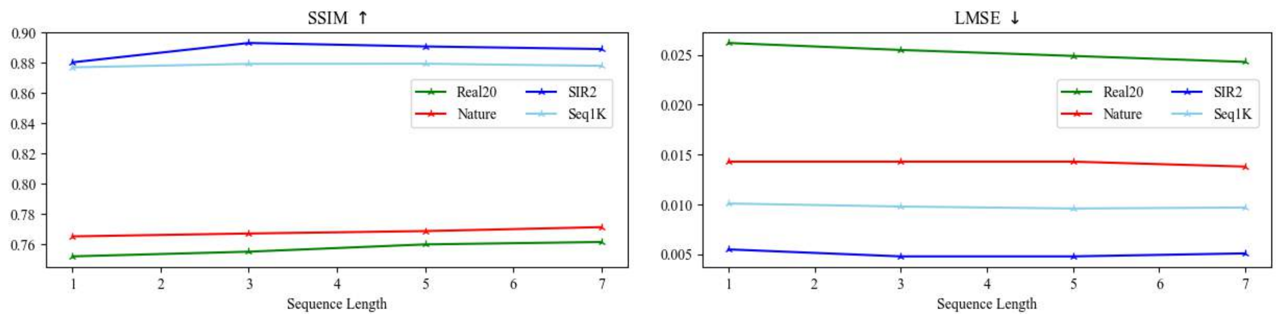

Figure 1. Effect of image sequence length.

	Real20		Nature		SIR^2		Seq1K	
	SSIM↑	LMSE↓	SSIM↑	LMSE↓	SSIM↑	LMSE↓	SSIM↑	LMSE↓
Syn.	0.7698	0.0231	0.7672	0.0140	0.8959	**0.0043**	0.8804	0.0096
Seq1K	0.7924	0.0212	0.7911	0.0136	0.8850	0.0053	0.8910	0.0084
Syn. + Seq1K	**0.8048**	**0.0200**	**0.7941**	**0.0129**	**0.8962**	0.0046	**0.9020**	**0.0072**

Table 1. Impact of synthetic (Syn.) and our real (Seq1K) datatsets on the performance of our model on four real-world datasets.

2. Effect of Input Image Sequence Length

To explore the effect of the length of input image sequences to our teacher network, we experiment with training our model using different image sequence lengths and test the resulting models on four evaluation datasets. Figure 1 shows that compared with training using single images, training with image sequences can lead to performance improvements. In addition, the models trained on longer image sequences better performance. However, we also note that, using a image sequence length of greater than 3 does not lead to further performance improvement. This is perhaps because a sequence of 3 multi-view images already contains sufficient motion information for our model to learn reflection dynamics for removing reflections well. Thus, we choose to use sequences of 3 images in our experiments.

3. Impact of the Synthetic and Real datasets

To analyze the influence of our synthetic and real data, we experiment with training our on each or both of the two datasets. The results are shown in Table 1. Compared with using either the synthetic data or the real data only, training on the combined data gains the best performances on three real-world datasets (*i.e.*, Real20, Nature and Seq1K). This suggests both datasets are important to the superior performance of our method.

4. Examples of Our Synthetic and Real Datasets

We show some examples from our synthetic and real datasets in Figure 2 and Figure 3. Figure 4 show some examples of reflection images in our real dataset. It can be seen that our real dataset covers a wide range of scenes, both indoor and outdoor, and different lighting conditions, both light and dark.

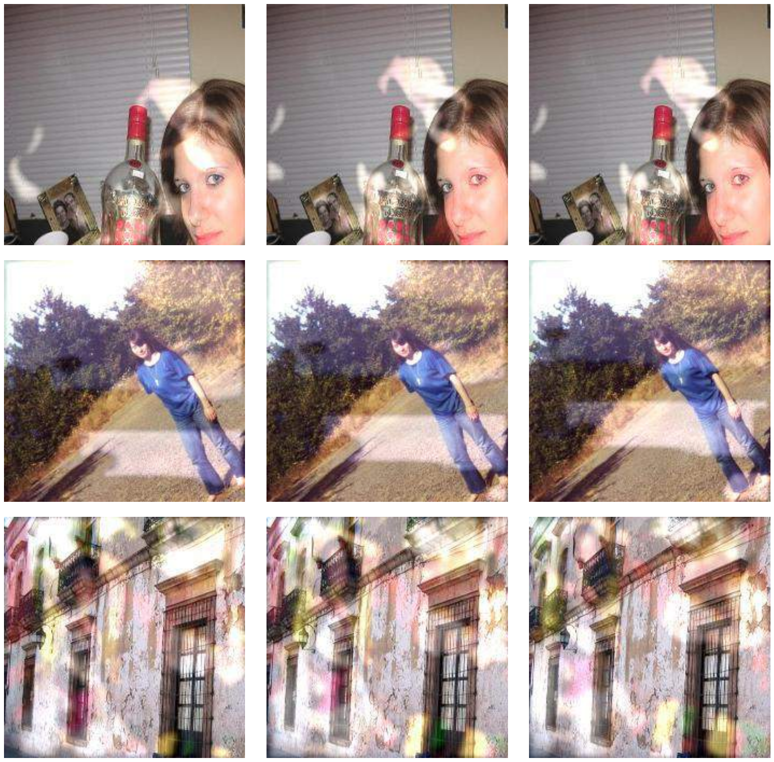

Figure 2. Examples of synthetic reflection image sequences.

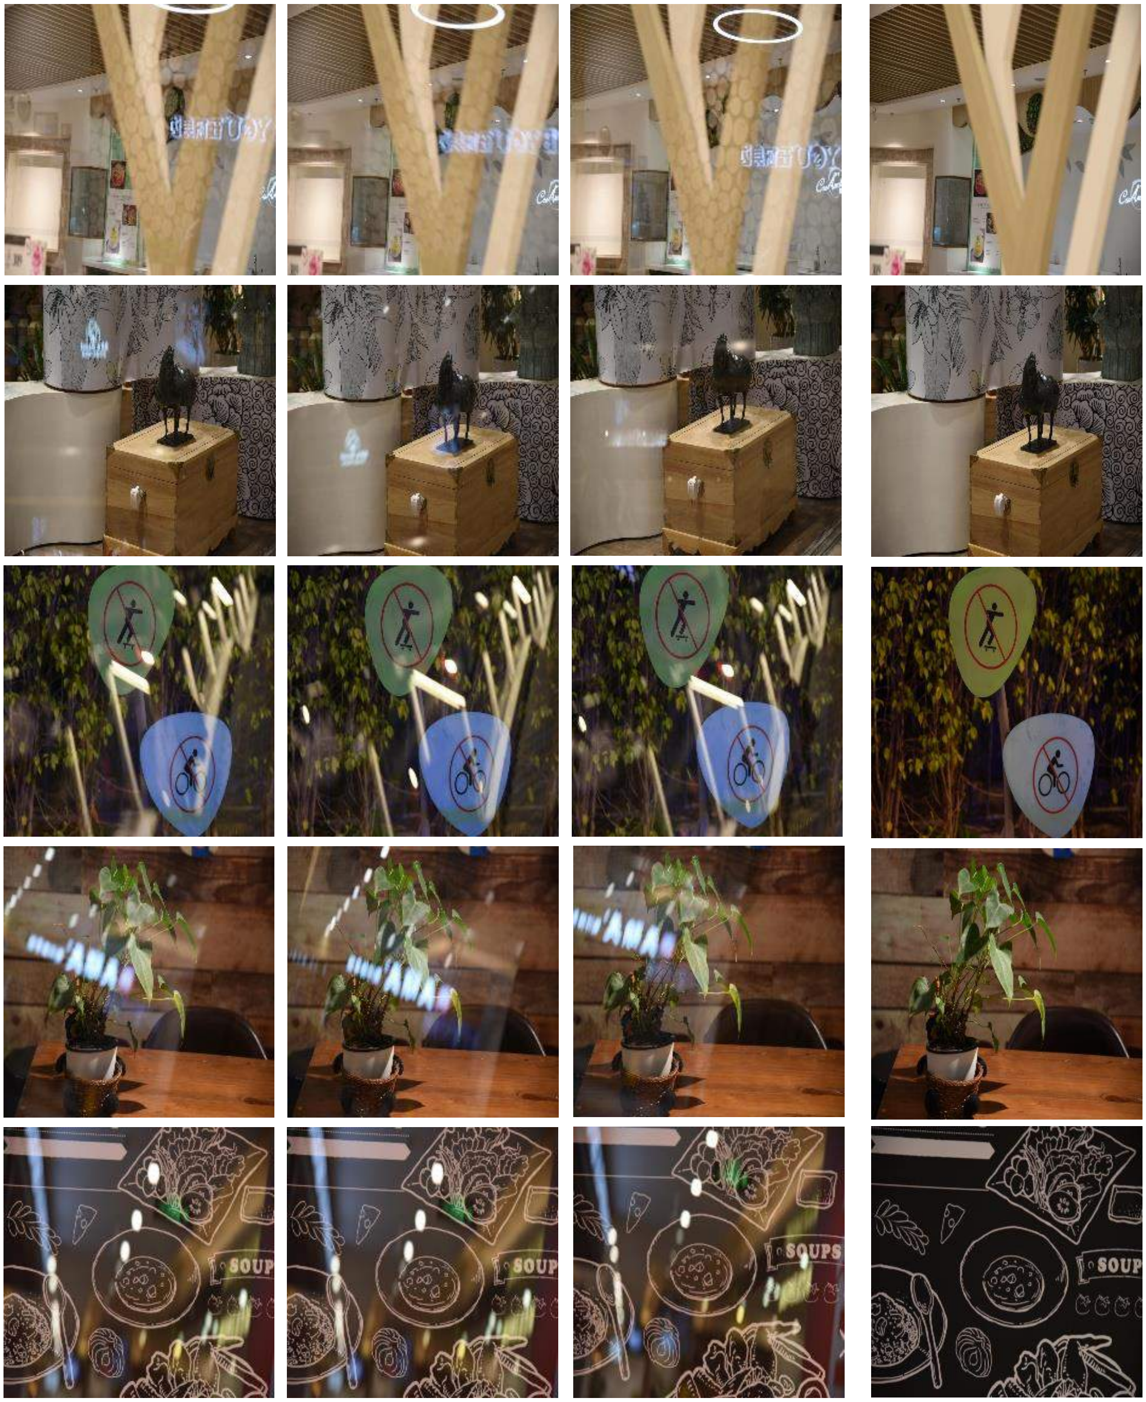

Figure 3. Some image sequences from our real dataset (left) along with the ground truth transmission image of the center image in each sequence (right).

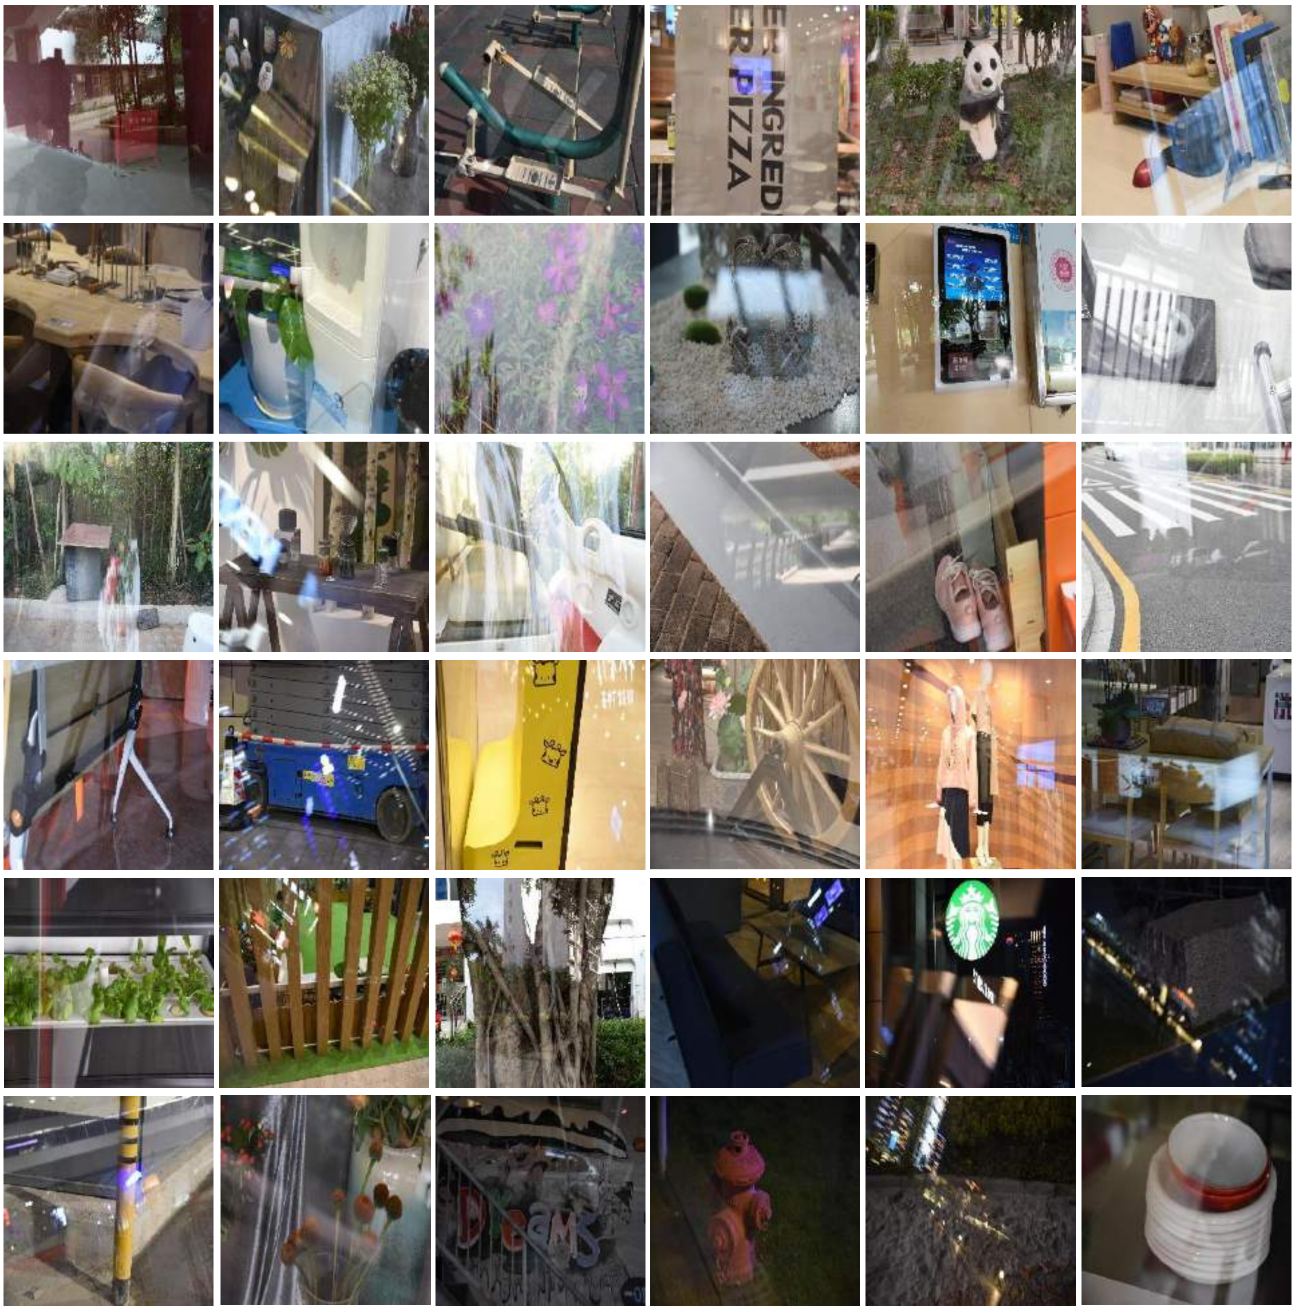

Figure 4. Examples of collected reflection images in our read dataset.

5. More visual results

Figure 5. Visual comparison of our method against prior methods on real-world images.

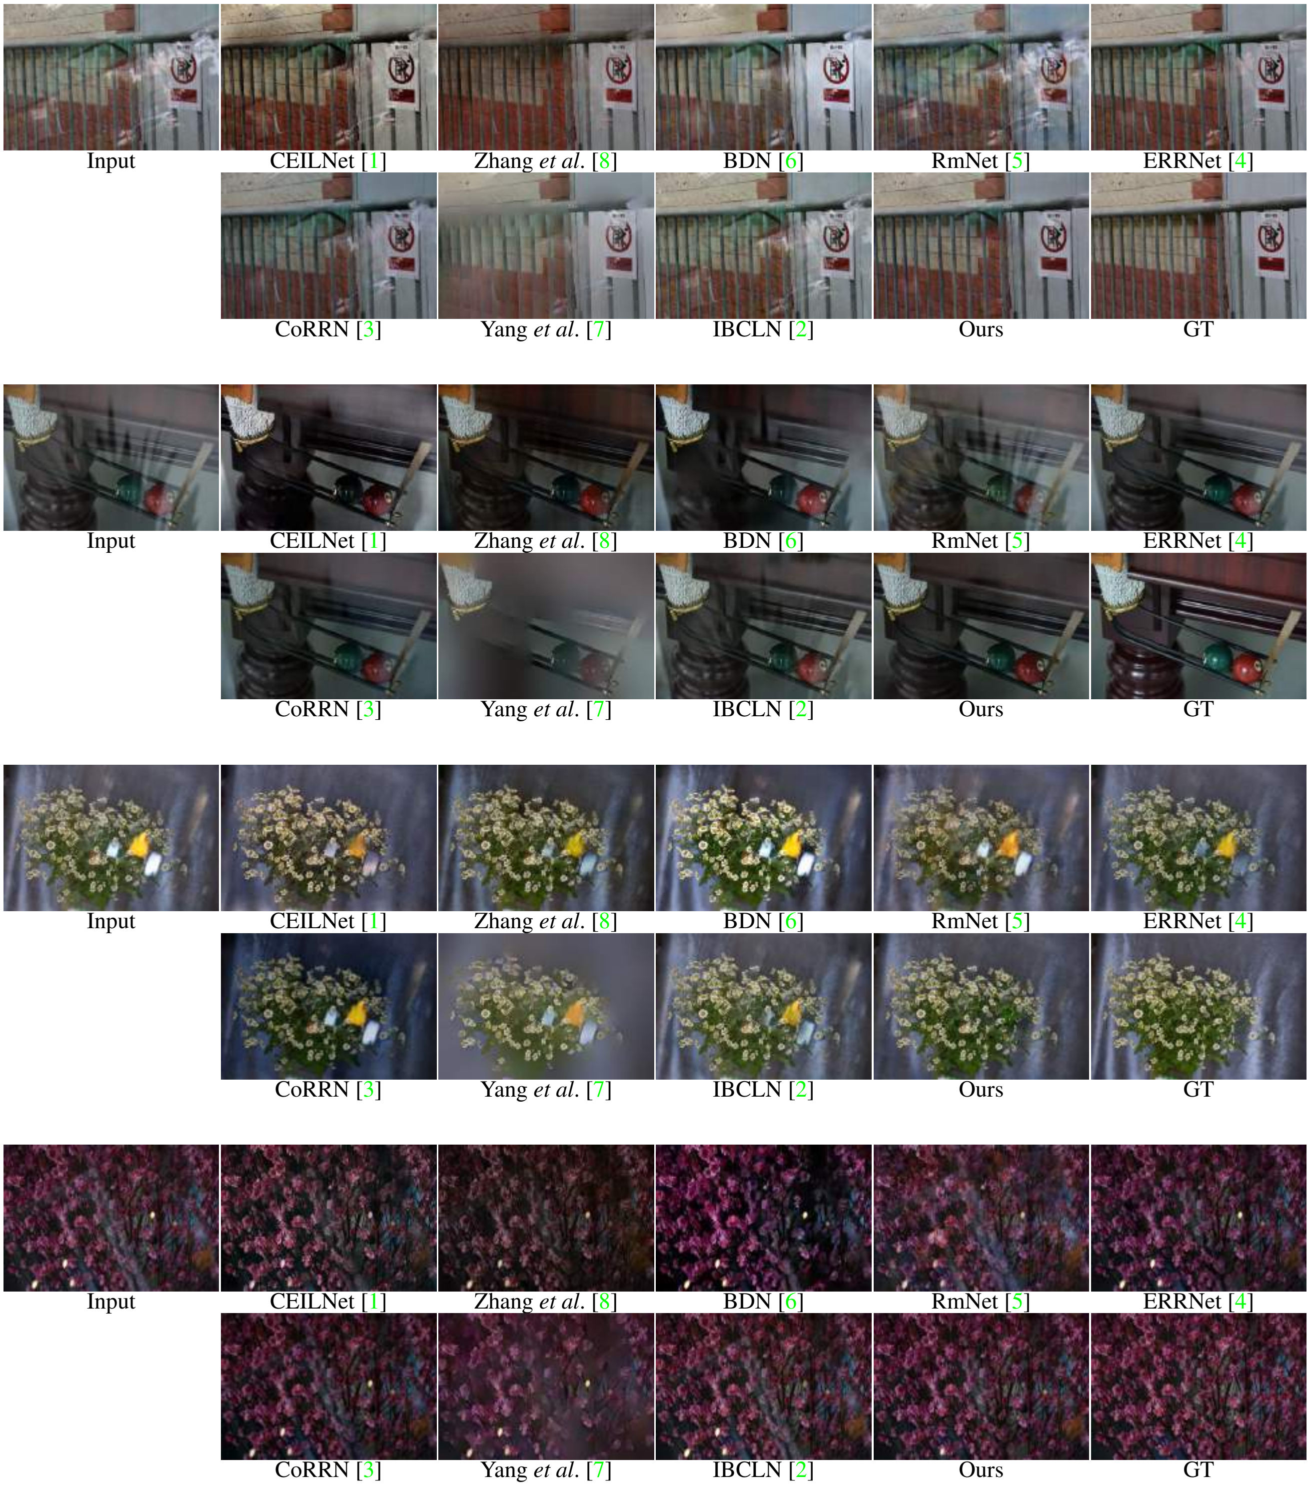

Figure 6. Visual comparison of our method against prior methods on real-world images.

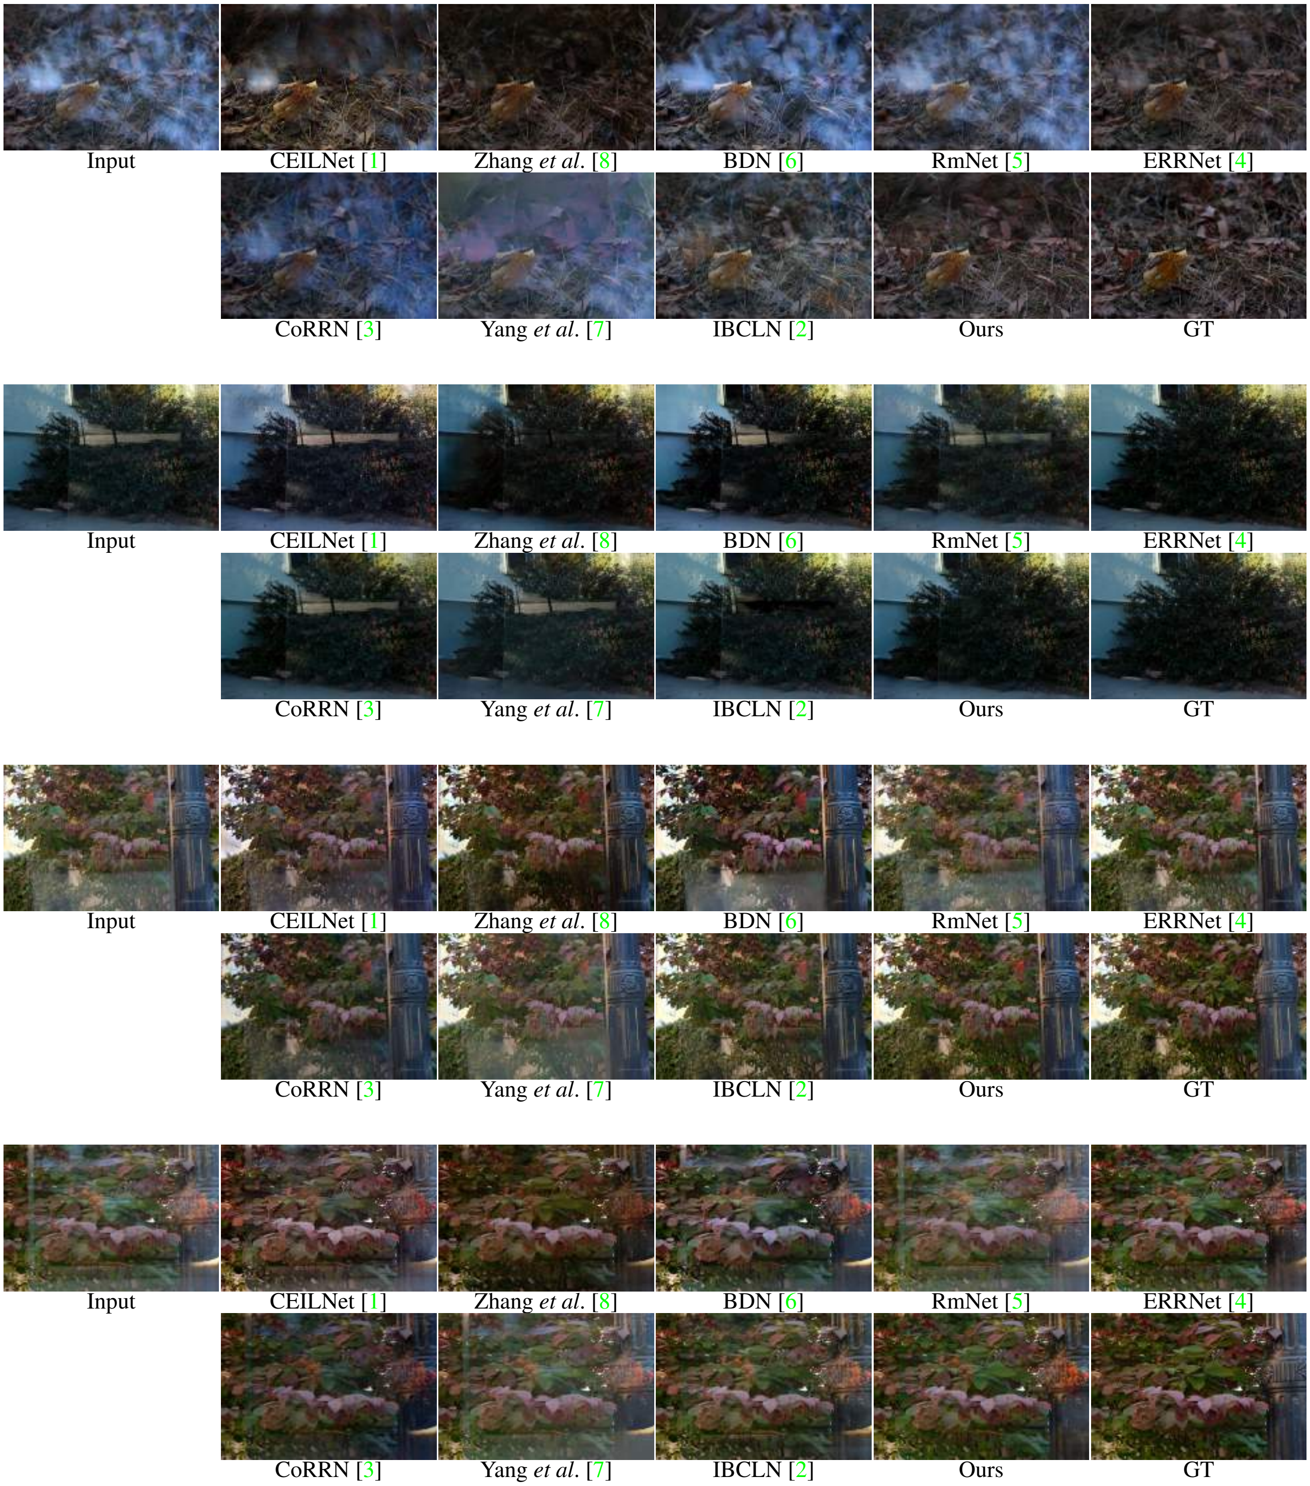

Figure 7. Visual comparison of our method against prior methods on real-world images.

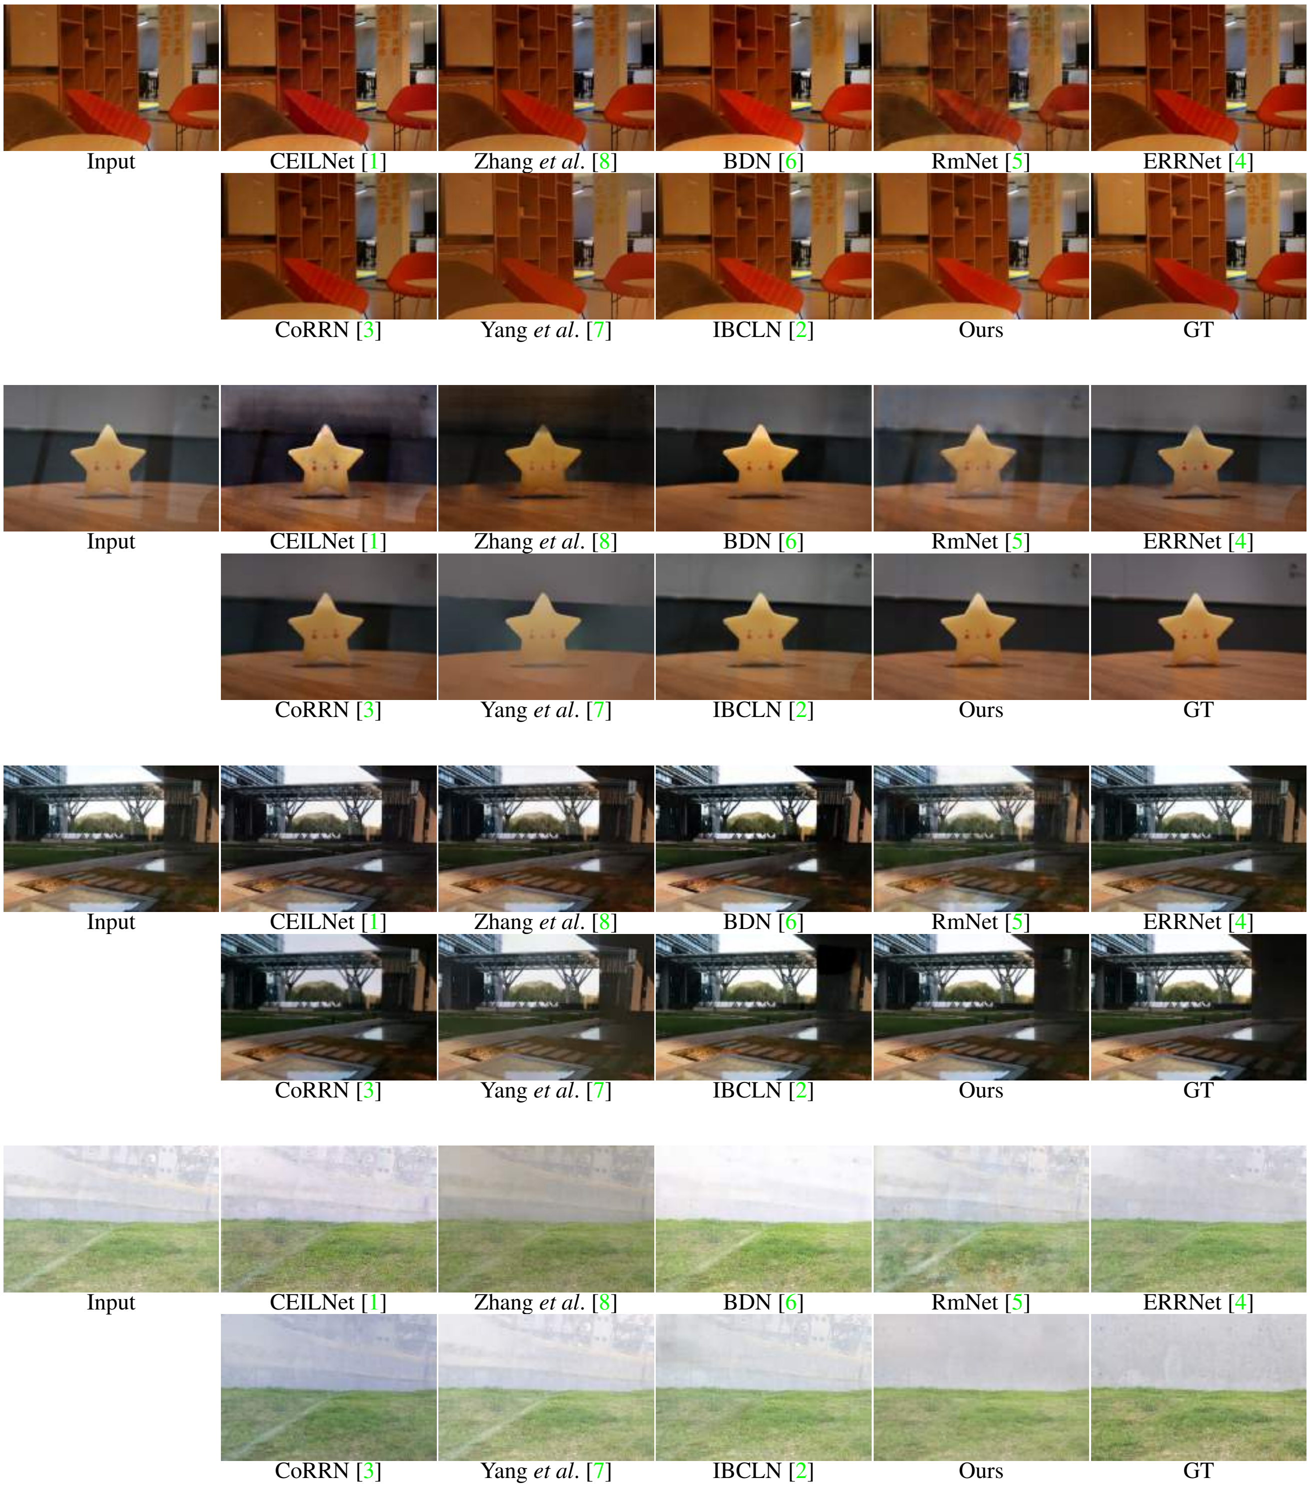

Figure 8. Visual comparison of our method against prior methods on real-world images.

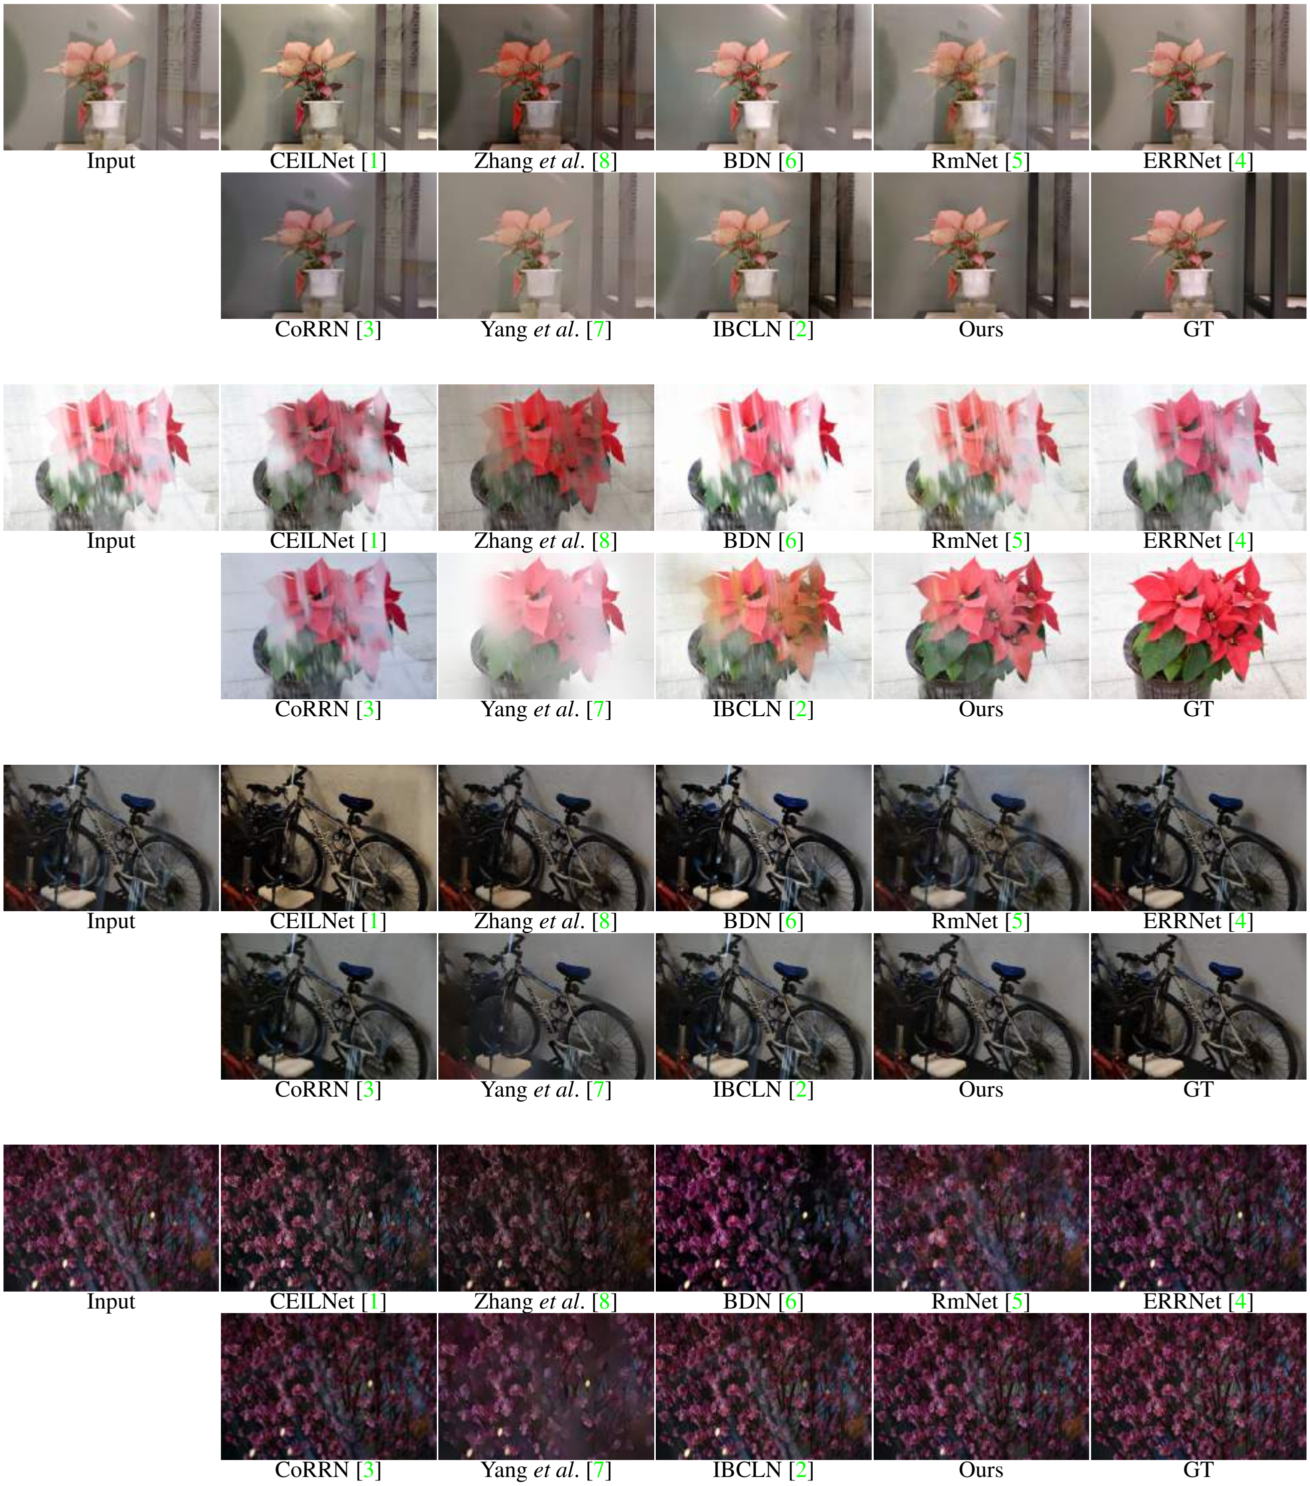

Figure 9. Visual comparison of our method against prior methods on real-world images.

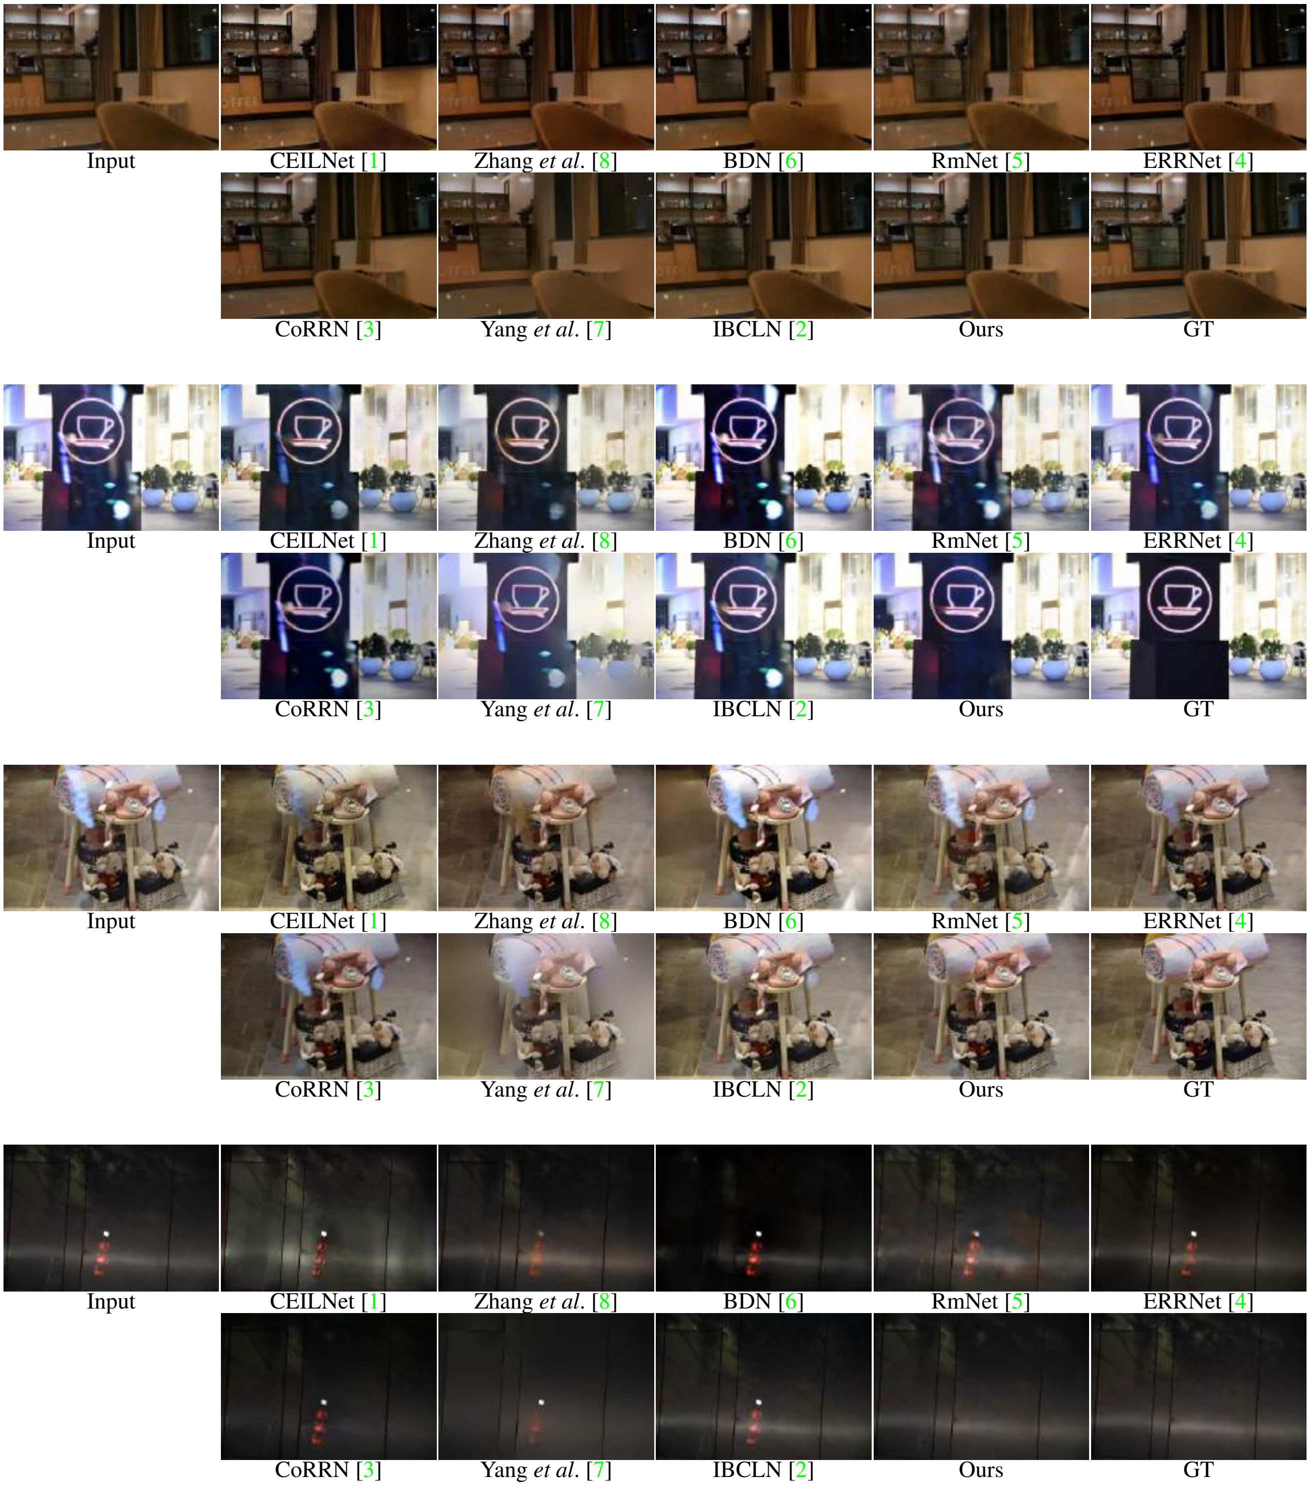

Figure 10. Visual comparison of our method against prior methods on real-world images.

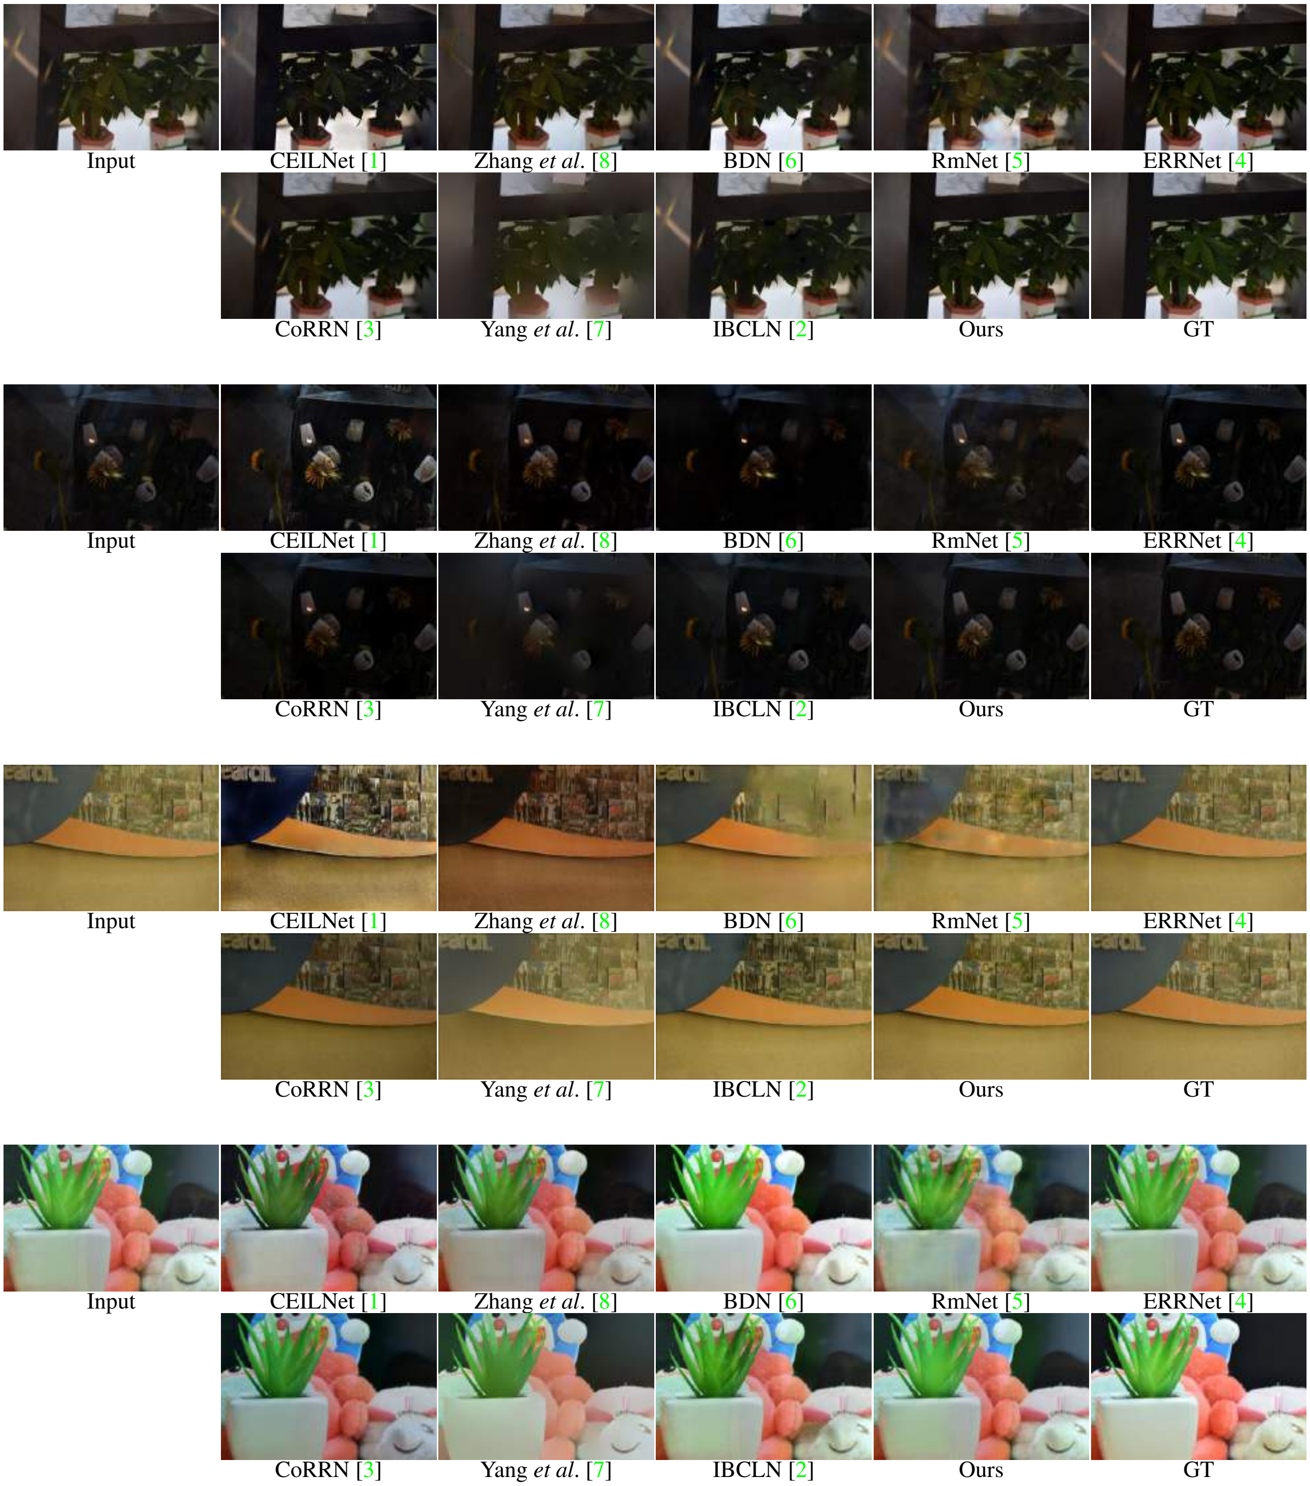

Figure 11. Visual comparison of our method against prior methods on real-world images.

References

[1] Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, and David Wipf. A generic deep architecture for single image reflection removal and image smoothing. In *ICCV*, 2017. 6, 7, 8, 9, 10, 11, 12

[2] Chao Li, Yixiao Yang, Kun He, Stephen Lin, and John E Hopcroft. Single image reflection removal through cascaded refinement. In *CVPR*, 2020. 6, 7, 8, 9, 10, 11, 12

[3] Renjie Wan, Boxin Shi, Haoliang Li, Ling-Yu Duan, Ah-Hwee Tan, and Alex Kot Chichung. Corrn: Cooperative reflection removal network. *PAMI*, 2019. 6, 7, 8, 9, 10, 11, 12

[4] Kaixuan Wei, Jiaolong Yang, Ying Fu, David Wipf, and Hua Huang. Single image reflection removal exploiting misaligned training data and network enhancements. In *CVPR*, 2019. 6, 7, 8, 9, 10, 11, 12

[5] Qiang Wen, Yinjie Tan, Jing Qin, Wenxi Liu, Guoqiang Han, and Shengfeng He. Single image reflection removal beyond linearity. In *CVPR*, 2019. 6, 7, 8, 9, 10, 11, 12

[6] Jie Yang, Dong Gong, Lingqiao Liu, and Qinfeng Shi. Seeing deeply and bidirectionally: A deep learning approach for single image reflection removal. In *ECCV*, 2018. 6, 7, 8, 9, 10, 11, 12

[7] Yang Yang, Wenye Ma, Yin Zheng, Jian-Feng Cai, and Weiyu Xu. Fast single image reflection suppression via convex optimization. In *CVPR*, 2019. 6, 7, 8, 9, 10, 11, 12

[8] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *CVPR*, 2018. 6, 7, 8, 9, 10, 11, 12